Expanded Tele-Health Platform for Android

Jochem van Gaalen
Phillip Azar
Adarsh Mani
Quan Peng
Daniel Aranki, Ed.
Ruzena Bajcsy, Ed.

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2015-86
http://www.eecs.berkeley.edu/Pubs/TechRpts/2015/EECS-2015-86.html

May 14, 2015

Acknowledgement

I would like to acknowledge the help and guidance of my mentors Dr. Ruzena Bajcsy and Daniel Aranki (Ph.D candidate).

University of California, Berkeley College of Engineering

MASTER OF ENGINEERING - SPRING 2015

Electrical Engineering and Computer Sciences

Robotics and Embedded Software

EXPANDED TELEHEALTH PLATFORM FOR ANDROID

JOCHEM VAN GAALEN

This **Masters Project Paper** fulfills the Master of Engineering degree requirement.

Approved by:

1. Capstone Project Advisor:

Signature: ___________________________ Date _____________

Print Name/Department: RUZENA BAJCSY/EECS

2. Faculty Committee Member #2:

Signature: ___________________________ Date _____________

Print Name/Department: ALI JAVEY/EECS

2015

Expanded Tele-Health Platform for Android

FINAL CAPSTONE REPORT

PHILLIP AZAR, ADARSH MANI, JOCHEM VAN GAALEN, QUAN PENG

Table of Contents

1. Introduction

In 2010, $17 billion dollars was spent on patient readmissions in Medicare. (Centers for Medicare & Medicaid Services: 2012). The most preventable readmission factors - and likely the most severe - are improper administration of therapy and medication. To combat this, we are developing a lean, smartphone-based telemedical system that not only provides personal health statistics to the patient, but relays this information to the doctor allowing for a more tailored approach to medicine. By leveraging the latest open source health monitoring algorithms and libraries, we have created a reliable, low-cost and easily implementable solution that has the potential to simultaneously improve patient care and dramatically reduce healthcare expenditure.

The team, advised by Dr. Ruzena Bajcsy and led by Ph.D. student Daniel Aranki, consisted of several members, including four Master of Engineering students (Phillip Azar, Adarsh Mani, Quan Peng, and Jochem van Gaalen) specializing in Robotics & Embedded Software concentration in the EECS department, and engineers Arjun Chopra, Priyanka Nigam, Sneha Sankavaram, Maya Reddy and Qiyin Wu.

The aim of this project is to develop and test a highly configurable, open-sourced Android powered framework that will enable doctors and hospitals to remotely monitor outpatients with chronic health conditions and determine their risk of re-admittance. This platform will transform the ubiquitous smartphone into the backbone of a tele monitoring system which uses various sensors to track different vital signs. These sensors can either be off-the-shelf products (e.g. temperature sensors) or sensors embedded in the smartphone itself (such as gyroscopes, accelerometers, cameras etc.). Once the data is extracted, the smartphone will be used to relay the information, in a fault-tolerant manner[1], to servers where various diagnostic algorithms are run to calculate re-admittance risk factor and other recuperation metrics. Security and privacy procedures will be embedded into these data structures prior to storage.

The team's specific contribution to this project is to develop:

1. Fault-tolerant communication between various off-the-shelf vital sign extractors and the smartphone via Bluetooth.

---

[1] Fault-tolerance can be thought of as a personal guarantee to the user that the process will be handled despite the presence of faults such as unexpected power downs, disconnection etc.

2. Novel vital sign extraction algorithms for heart rate and blood pressure using the smartphone's camera to serve as a proof of concept to demonstrate the fault-tolerant communication and provide a simpler way to conduct health and vital monitoring.

Over the course of the academic year, an application in Android was built and tested through laboratory validation of individual components, and a pilot study conducted near the end of the project which combined these. The results of this work is outlined in this report, along with challenges faced and potential next steps for this project.

This project was also analyzed from a business perspective, with potential competitors highlighted, and a thorough analysis of the industry and market trends to aid in identifying potential business opportunities.

2. Market Trends Analysis

2.1.1 Motivation & Outline

This section of the paper is dedicated to analyzing the telehealth industry and deriving a potential business strategy for a startup centered about our capstone project. It begins by discussing the broad industry trends that show which direction would yield the best go-to-market results. Next, the importance of adopting a well-defined strategy in terms of maximizing business potential is highlighted. From analyzing the industry as a whole in terms of its value chain, the most optimal place within this market a business should occupy will be determined. This is done by identifying the full chain and finding the strongest link. This is followed by an analysis of Porter's Five Forces, which is used to understand strengths and weaknesses in the chosen position within the value chain. Knowing these strengths and weaknesses is a key factor in shaping the business plan. The conclusions drawn from these forces will be presented along with suggestions for ameliorating any weaknesses identified.

After extensive analysis using these models, the section of the market that was identified to be targeted is the "mediator" segment - a segment whose products would form the platform that will serve as common gateway for devices measuring vital signs to communicate with health analytics services and vice versa. The platform the team is developing will allow for a much easier facilitation of data between entities specializing in either the analysis or measurement of health data. It essentially abstracts away the communication layer of the entire telehealth chain into an easy to use, robust and secure service. The strategy adopted to conquer this segment is to pursue an open-source double licensing model wherein royalty will be collected at a pre-determined rate only when the service or code is used for commercial purposes. Personal use cases will be free of charge.

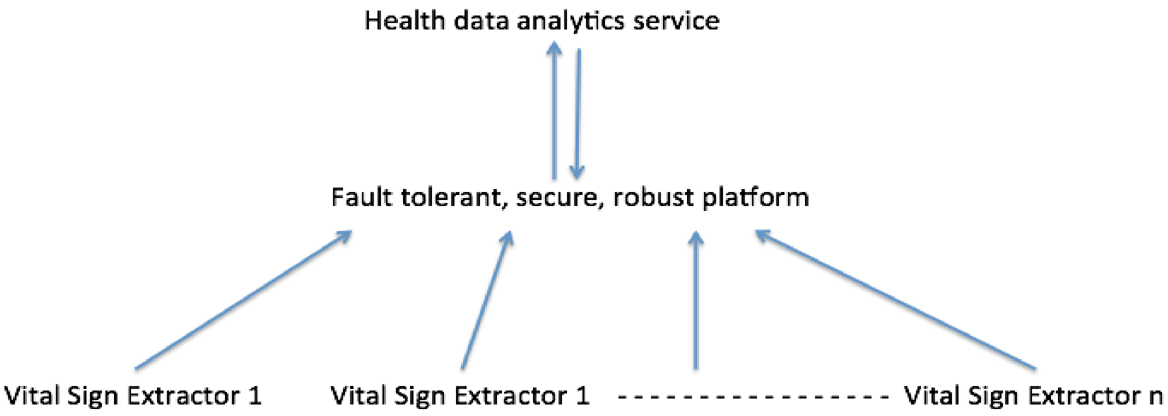

Figure 1: Telehealth Platform Hierarchy

We live in a highly interconnected world. Everything from common household appliances, such as refrigerators, to complex control systems, such as those found in vehicles are connected to information relaying services designed to expand their ordinary functionality. Our lives are becoming more and more dependent on being "connected," and as a result, many major social, technological, and economic trends in the last decade have been fueled by connectivity. Among these trends is the push for telehealth, which since the 1990s, has been a principal force in providing healthcare to underserved communities. (Miller 2007: 133-141) By leveraging major social trends towards connected privacy, availability, and transparency, it is possible to flesh out a telehealth strategy that addresses key impacts and vantage points associated with what will be collectively known as the "connectivity trends." These connectivity trends are driven by a new, socially aware consumer that can greatly influence a product's adoption rate and popularity.

To provide context for these consumers, it is important to examine contemporary cases of social awareness with respect to connectivity. In recent years, the online health industry has been abuzz with issues of privacy, data storage, and patient protection. So much so, in fact, that entire companies have been spawned out of the necessity for these core three attributes (e.g. Veeva Systems.) Chief among these concerns is that of patient protection. Where as many medical devices remain "analog" with respect to connectivity (i.e. no internet/wireless access,) manufacturers are pushing to create connected equivalents of these products. With this comes the potential for new features, but also the introduction of many risks. In a more extreme (and relevant) example, data from an insulin pump that now communicates via Bluetooth low energy may find itself in the hands of third party advertisers. (Hall & McGraw 2014: 216-221) While this data is protected under the Health Insurance Portability and Accountability Act (HIPAA) passed in 1996, it is very likely that the patient will be unaware of what is occurring, since they are often surrounded by Bluetooth enabled devices with internet access. This creates a very real, and potentially dangerous situation: data can be intercepted and collected from medical devices, which can then be used for almost any purpose. And in effect, what you find is similar in many ways with the motivation behind the Digital Millennium Copyright Act (DMCA,) in which illegally obtained data is so abundant that liability is no longer held against companies who store that data.

From a strategy perspective, the contemporary social awareness of patient privacy, data storage, and protection significantly detracts from attempts at commercialization. The FDA Center for Devices and Radiological Health (CDRH) is the major regulatory arm that governs the development, implementation, and evaluation of products with telehealth features. Since the FDA is a member of the federal government, policymakers, who in turn are influenced by those who they represent (and in many cases, special interest groups,) have the ability to directly impact the regulation of telehealth systems. Negative social acknowledgement can therefore directly harm the commercialization of a new telehealth product by influencing the implementation of stricter regulation. So the question remains, how can telehealth products hope to compete against an increasingly aware consumer base?

It is not feasible to simply promote ignorance of a very real and tangible problem: the digital age has long since begun, and next to nothing is primarily stored physically. If you cannot simply "wish" the knowledge away, then the next best thing is to approach the problem with the contrapositive: education. And this kind of education, one that targets the socially aware consumers (and in this case, patients,) is not strictly limited to the benefits of a product. In fact, it is critical to convey both sides of the technology. Net neutrality is one modern example where conveying only the benefits of a technology is more subversive than it is helpful. Major internet service providers attempted to only convey the benefits of discarding net neutrality in an attempt to mask its true purpose as a means for generating additional revenue. This had the effect of creating a nearly unanimous voice against the proposition. In conveying both the benefits and downsides to telehealth, we can best leverage the trend of social awareness in connected health care, and in turn, create a more positive public image.

The importance of influencing the socially aware consumer cannot be understated: coordinating a strategy of education rather than marketing can mean the difference between revolutionizing and downplaying the entire telehealth industry. With the advent of free online educational resources, it is simply not good enough to assume that a market segment will be ignorant about a product or feature. A good strategy will leverage this education and in some ways seek to exploit it, feeding a growing demand for high tech while at the same time appealing to the scientific side of a market segment. There are, however, many more sides to this "good" strategy that cannot be addressed by analyzing the consumer alone.

2.2.1 Strategy: What it is and why it is required

Strategy is both the art and science involved in the formulation and evaluation of key-decisions that will give an entity an edge over its competitors, attain a sustainable advantage and help achieve its long-term objectives. Delineating a strategy also helps the organization formalize objectives, and hence give it concrete direction. It helps one decide how to position a product and what decisions to make to maximize growth and profitability, as it is just as important when trying to enter a saturated market as it is when trying to enter a nascent one. It is just as useful when trying to consolidate one's lead in a market as when one is trying to establish it.

The telehealth industry is still relatively nascent, but is rapidly expanding. Industry revenue was projected to grow $320.2 million in the five years leading to 2014, including revenue growth of 23.1% in 2014. From 2014 through 2019, the industry stands to reap the benefits of demographic, structural and legal factors. These include a growing aging population, advances in telecommunication and wearable technology, and statutes like the Affordable Care Act. As a result, industry revenue is expected to increase at an annualized 49.7% to $2.4 billion in the five years to 2019 (Morea 2014:9).

Although the top two companies, GlobalMed and InTouch Technologies, occupy 22% of the market share (Morea, 2014), the industry itself is rapidly expanding, and thus there is ample opportunity for this technology to enter the market and seize a sizable share.

However, given that the project is a potential spin off from an academic endeavor, and doesn't have (at least presently) the financial clout or business savviness of companies already in the market, adopting the right strategy becomes crucial to conquer and succeed in this market. Where the team temporarily lacks in business authority, it excels in technical skill and prowess. While it is acknowledged that this will have to change in the future, the team believes it can use this technical superiority to get a foothold in market.

2.3.1 The Strategy

As mentioned in the previous section, the team plans to leverage its technical prowess to its advantage. This involves building a robust easy to use, easy to integrate platform, and then treading the open-source route. In the current market, existing solutions are mostly "walled-gardens" with players providing proprietary solutions to all fields in telehealth (vital-sign extraction, collection, logging, analysis and result distribution). However, with the impending explosion of the internet of things, several third-party developers of medical hardware used for constantly monitoring vital sign data are expected (Lars 2014). The number of connected devices is expected to reach 50 billion mark by 2020 (Swan 2012). This data emphasizes a need for an open platform to handle the software/communication part of the system. It is this opportunity that the strategy is designed to seize.

2.4.1 Industry Value Chain Analysis

The first requirement one must consider when creating a new business is how it should be positioned within the industry. For example, there are many levels involved in delivering medical devices to hospitals. A company wanting to produce a device must first buy supplies, and then usually contract out the labor to a manufacturer. Once the devices have been manufactured, the company will sell these medical devices to a distributor who carries the device in a catalog. Hospitals partnered with these re-sellers will then purchase these devices, process them, and finally hand them over to doctors who will use them as a part of the care they provide.

We can see that the value chain is long, and in many cases, not quite linear. Various market forces also make the chain asymmetric in that the value that is realized in each part of the chain is not evenly distributed. Many markets often have a particularly 'fat', or highly profitable part of the chain while the other sections fight for the leftovers. An example of this can be seen in the technology industry. In delivering cell phones to consumers, Apple, the company that designs and markets its products, clearly makes the most money, even though there are many other businesses involved in the value chain such as the phone manufacturers (e.g. Foxconn), distributors (e.g. FedEx) and carriers (e.g. AT&T). Unlike the consumer technology industry which arguably has started to reach maturation, telehealth is still a fledgling industry, which is still being shaped by early startups, new government regulations and many new technological innovations. Despite this, the telehealth industry value chain is no different in that the 'fat' part of the chain can clearly be identified.

From the IBISWorld report on telehealth, the industry value chain is shaped with several 'demand' players such as data processing and hosting services, hospitals, physical therapists and health & welfare funds to name a few (Morea 2014:14) . On the other side of the chain, there are several 'supply' players providing infrastructure within which telehealth services can operate. These include businesses such as mobile carriers, software publishers and hardware manufacturers.

Because of the wide variety of players in the industry, both supplying and pulling technologies to and from telehealth services, a simplified linear model was developed to illustrate the supply chain. This is shown below in Figure 2.

Figure 2: Industry Value Chain

From here, we see that the chain is longer than one might expect in that there are several intermediate steps involved with getting a device into a hospital. It is important to note that this value chain is a simplification of the actual chain, which is not linear and has many more steps involved within each major stop in the chain presented.

Due to the established nature of similar industries requiring the same suppliers and hardware manufacturers as the medical devices industry, suppliers and hardware manufacturers are well developed and highly competitive (Kahn 2015:24). Thus, the fat part of the value chain is located with the medical device designers. A good example of the power of this part of the chain are medical devices like Medtronic and Siemens AG - companies whose revenues exceed $17 billion USD**,** which is where we would like to position a potential startup. It is important to note however that there remain several large pitfalls in which a business situated as a device designer could become wildly unprofitable. Many of these have to do with regulatory approvals (Morea 2014:21) such as meeting FDA requirements (Medina, Kremer & Wysk 2012). However, once met, these businesses will be well situated within the chain. The team's particular business would circumvent many of these pitfalls by taking on the 'mediator' role within the medical devices industry. As alluded to earlier, this would allow participating medical devices companies to safely gather data from patients and provide additional services derived from this.

By identifying the best location within the industry chain, we can now analyze which parties in the medical industry would have interests driving the adoption of new medical devices. These parties would be ideal candidates to which a startup business could offer services. Performing this analysis also has the added benefit of being able to predict the demand of the industry, which is a function of the size and strength of desire of its customers.

2.5.1 Potential Stakeholders

By its very nature, there are a plethora of potential stakeholders with key interests in the success of this platform. First, insurance companies would have an interest in the rise of tele-health monitoring systems, as having the ability to better assess the risk of patients could drastically change the way that insurance is calculated. From this, it follows that insurance payers and patients would see substantial benefits in the way that medicine is applied in that increased information about the patient will allow for more accurate diagnoses and more targeted treatments. This will ultimately help in lengthening and improving the quality of life.

Furthermore, the ability to take measurements of a patient's vital signs opens the door for alternative options to physically visiting the doctor – a practice continued for centuries. This is an enormous market currently valued at approximately $1.151 billion USD and growing at approximately 2.2% yearly (IBISWorld Report: Number of Physician Visits 2014:1) which presents a huge financial opportunity for the telehealth medicine. While certainly not a complete substitute, the telemedicine industry could replace large swathes of routine checkups and minor cases. This could serve as a boon not only for businesses looking to save on costs associated with granting time off to see a doctor but also for understaffed hospitals to relieve labor-related costs.

Governments and regulatory bodies such as the World Health Organization (WHO) would benefit from enhanced data collection in order to present new legislation or seek out cost cutting opportunities in relation to providing health care, and to help model and eventually mitigate the effects of contagious new diseases. With recent legislative moves seeking to expand health coverage in the US such as the Affordable Care Act, an increasing amount of financial support will flow into the telehealth industry contributing to the estimated 49.7% growth in industry revenue (Morea 2014:5) as a reaction to making care affordable to a wider swath of the population.

From this analysis, we can see that there are many large and powerful stakeholders with keen interests to either increase profits, or reduce costs with the aid of telehealth technology. This further validates the team's judgment in the Porter's power of buyer's analysis. With so many parties interested, and relatively weak internal industry rivalry, starting a business in this sector of the industry should have the highest chance of success.

Now that the stakeholders have been identified, it is crucial to understand the role that the team's technology would play in commercial applications. Investors in a startup are investing primarily in the technology, and the team, so without a clear vision of how the technology will function, it would become impossible to market it correctly or draw up a business model.

2.6.1 Relation between Technology and Commercial Product

The ability to monitor vital sign information in real time has direct applications to commercial technology. The team aims at providing a common platform to developers with some sample vital sign extraction methods embedded with the idea that other companies build on top of this platform. This platform will enable developers to work on producing technology more focused on data extraction and/or analysis without having to spend inordinate effort establishing fault tolerant protocols, server security, and data privacy. The idea here is to position the platform similarly to how the Android operating system is marketed. Android is an open-source cellular operating platform that comes with several basic functionalities, but is meant as a platform where third-party developers can extend its functionality.

Different vital sign extraction techniques are analogous to new apps developed specifically for this platform. Revenue would be collected from licensing agreements between device makers seeking to utilize the platform for profitable ventures. Such a strategy would have a benefit of rapid growth due to its flexibility, but could have the potential downside of being slow to grow revenue. In all likelihood, due to the nature of the technology industry to rapidly grow in new sectors, this would be the best approach as opposed to a strategy that would involve keeping the platform closed and competing rather than collaborating with medical devices and technology companies.

With a defined position within the value chain, and an idea of how the technology in this capstone project will interact with the final product, we can now analyze the remainder of Porter's Five Forces to assess potential threats in the industry.

2.7.1. Porter's Five Forces

In his seminal article in the Harvard Business Review, Michael E. Porter outlines five major factors that should influence the approach one should espouse to be successful (Porter 2008). These are:

1. Bargaining power of buyers

2. Bargaining power of suppliers
3. Threat of new entrants
4. Threat of substitutes
5. Rivalry amongst existing competitors

Each of these forces will be analyzed in detail in how they pertain to our chosen industry in the following sections.

2.7.2 Bargaining power of Suppliers

The bargaining power of suppliers refers to the ability of business providing materials, data, IP etc. to a business further down the chain to ask for increased compensation or to switch to supplying a direct competitor. Due to the fact that this platform is largely based off the heavily subsidized, and highly competitive mobile phone market, the suppliers in this regard have little power to demand higher prices were the telemedicine industry to grow and increase their demand. Furthermore, the ubiquity of Android-based phones (International Data Corporation 2014) in the consumer market further weakens dependence on suppliers as our platform does not require a proprietary phones. Rather, a patient can just use a personal phone to use the monitoring platform.

2.7.3 Bargaining power of Buyers

The bargaining power of buyers refers to the ability of customers of the business to demand extra features, lower prices, or switch to a competitor. Due to the suggested licensing strategy of this product, the buyers in this case would be device manufacturers and potentially research laboratories and regulatory bodies looking to do research on the aggregate dataset of patient data. It is without doubt that the sale or licensing of data would come with some sort of public resistance, but if done in an identity sensitive way (and this is already ingrained into the platform itself from its inception) this could open the door to many licensing opportunities. Since there are few alternatives to this system, the power of buyers to simply switch to another platform would be weak.

2.7.4 Threat of New Entrants

New entrants bring with them fresh ideas and a tenacity ideas to gain a share of the market from existing players. Ease of entrance, characterized by barriers to entry, will encourage more new comers especially if the industry is nascent and has potential for being lucrative. This precisely characterizes the telehealth market. We live in a connected world and with the Internet of Things verging on explosion, telehealth is one of the many industries that will receive a significant boost in market size (Lars, 2014.)

In such a scenario, it is anticipated that many players will prop up to get a piece of the pie. It is, therefore, of utmost importance to increase the barriers to entry to dissuade potential entrants. The team has done this by increasing barriers to entry in the following ways:

1. As the team is building a platform that will be leveraged by third-party manufacturers of health monitoring devices, the platform has to be robust, complete and easily joinable. Developing a solution that offers all three features requires a lot of developmental effort and time, which in turn raises the barrier of entry.
2. Once adopted, it'll be very difficult for our customers/partners to switch to a new platform.
3. The open source model works in the team's favor. Open source provides for collaboration and cooperation as opposed to competition. As the platform gains more adoption, a larger community would organically grow, further fuelling adoption.

By raising the barrier to entry for new entrants, the team has ensured a sustainable "economic moat" can be built.

2.7.5 Rivalry

Rivalry is greatest in an industry if competitors are roughly of equal size, industry growth is slow or if exit barriers are high (Porter, 2008). It can become particularly entrenched if the rivals are competing on price. Size and financial might of the competitors play a major role in determining who comes out victorious in market share wars. The current major players in the telehealth market are GlobalMed and InTouch Technologies who together own 22% of the market (Morea 2014:25). However, the industry itself is rapidly growing and hence there is abundant opportunity for entry.

In the telehealth industry, the basis of competition is primarily based on product quality, which is determined by factors like product functionality, features, speed and ease of use. Other media of competition include breadth of services, marketing, customer services and training offered to clients. (Morea 2014:24)

Given this, the biggest threat for the team comes from the threat of rivals. Being a potential startup spun off from an academic endeavor, the team is definitely at a disadvantage in terms of the amount of money that can be spent on marketing and customer service.

However, the open-source strategy being adopted will help us overcome this hurdle. The team aims at revolutionizing the telehealth industry the way Linux revolutionized the mainframe operating system industry or the Android revolutionized smartphone operating system industry. By providing a free to use platform that can be leveraged by third party developers, the growth of a community of support and innovation is encouraged. This will be particularly successful because:

1) The proliferation of cheaper hardware and sensors is spurring the growth of independent original equipment manufacturers who can use the team's platform to build a complete solution.
2) With cheaper networking solutions and the growth of internet of things, several startups are developing devices each with their own proprietary communication protocols. Currently, there isn't a single platform that would allow these disparate devices to talk to each other. This reduces interoperability, which can stall growth of the industry a whole (Ghosh et al., 2011). As adoption increases, the open source platform the team is developing has the potential to allow for greater interoperability, thus further spurring the industry growth.
3) Existing companies have their proprietary platforms which can be licensed by third party developers at a hefty fee. By providing an open source solution, the team can win over more developers, thus gaining market share.

2.7.6 Threat of Substitutes

A substitute product is one that that offers similar benefits to customers of a company from other industries. The threat of substitutes describes the possibility that substitutes can replace the company's products. If the company is faced with a high-level threat of substitutes, it loses the control over the price that it can set to sell to customers, since customers can easily switch to its substitutes if the price is set too high. Therefore, it is important to analyze the threat of substitutes given it can affect profitability of the company and the industry. In this section, we will discuss the problem our technology tries to solve, what substitutes there are that solve the same problem, and what the level of threat the substitutes pose on our technology.

Health care is by far one of the most expensive industries in the United States. Readmissions lead to a huge hospital costs. According to Agency for Healthcare Research and Quality, there were 3.3 million readmissions in the United States, which cost hospitals about $41.3 billion USDs (Heins et al 2014). By developing a system that can reduce trips to the hospital, helping people get in touch with their own health, and mediate the link between patient and physician, we can cut down the waste in health care. Specifically, the problem that our technology tires to solve is to reduce hospital readmissions by reporting patients' vital signs (including heart rate, blood pressure, energy expenditure and body temperature) to physicians, who will in return give patients' health advice remotely to keep healthy.

Except our technology, traditional medical devices and modern technologies may achieve the same purpose and become our substitutes. Traditional medical devices are single functional devices such as thermometer, sphygmomanometer, EE monitors and etc. Patients can use those devices to record their vital signs by themselves and remotely consult physicians for feedback with recorded data.

Traditional medical devices have advantage on price over our product. They are relatively cheaper to Bluetooth-enabled medical devices required by our technology, which may lead patients to adopt traditional medical devices instead of our technology. Another weak point of commercializing our technology is switching cost. Since our technology is open-source, we should obey open source philosophy, which states that open-source software should be freely used, changed and shared (Open Source Initiative 2015:1). As a result, there is no reason to charge termination fees to patients when they decide switching to substitutes. However, our technology differentiates from these substitutes in three aspects which can stick patients to our product:

1. Instant data transmission.

   Our technology helps patients to transfer health data to physicians. If the patients use traditional medical devices to monitor their vital signs, they have to record the data and call or email physicians for advice. This procedure is both time consuming and boring. The patients may give up tracking their health conditions after several trials. With our technology, the health data gathered by Bluetooth-enabled devices will transmit to smartphones immediately and then transfer to physicians by cellular network. The instant health data transmission saves time for patients and physicians.

2. Easy access to health data.

   Our technology provides convenience for patients to check their vital signs record. Traditional medical devices are separate. To check vital signs history, patients need to check on individual devices. Our technology is based on smartphone. All measured vital signs will store on the phone. As a result, patients can check their health data anywhere anytime with their phone.

3. Expandable functionalities.

   Our technology is expendable because we built an open-source platform for telemedicine. Developers can write more applications using our platform. Therefore, in the future, our technology has potential to provide more useful functionalities as well as benefits to patients and physicians by the contributions of third-party developers.

In addition to traditional medical devices, modern technologies could also pose a threat to the platform we are developing. Following are two existing modern technologies serve as substitutes of our product.

1. Motorola - Moto 360™

Motorola has recently released a smartwatch called the Moto 360™. Other than being an accessory to a smartphone, it can give basic biomedical data such as skin temperature, heart rate and energy expenditure. This platform could be developed further to come closer to the product we envision, but this would require substantial modifications to the current product.

2. Microsoft – Kinect

Microsoft's Kinect system has cameras which can be used to detect heart rate, energy expenditure and even mood by sensing the changes in flood flow in the face (Zhang 2012:4-10). Microsoft is currently focused on integrating the Kinect system with its Xbox One gaming console. However, this does not stop third-party developers from using this system to provide in-house medical diagnostics such as respiratory surface motion tracking (Alnowami et al 2012), 4D thermal imaging system (Skala 2011 et al:407-416), and otoneurological examinations (Dolinay et al 2014).

With considerations about price, switching cost, and differentiators with respect to traditional medical devices and discussion about modern technologies, the threat of substitutes of our product is moderate.

2.8.1 Major Conclusions

In this report, it was concluded that the best part of the industry for a startup centered about the technology being developed in our capstone is as a mediator between medical device companies and data services. This is the fat part of the chain. A Porter's Five Forces analysis was also conducted, and it was determined that with small tweaks to the planned dual-licensing strategy, all of the forces can be reduced to weak or moderate in strength. It is important to note however that this favorable analysis in no way guarantees success. As will be discussed in papers to come, the largest hurdles lie not within the placement of the startup, but with issues like avoiding IP prosecution, gaining the necessary industry approvals, and building up a reputation for being protective about sensitive health information. These issues are not easy, and some will be addressed specifically in future papers.

3. Intellectual Property Strategy

3.1.1 Introduction.

This section of the report is dedicated to explaining the strategy the team would like to adopt with respect to intellectual property. While the final goal is to create an open source product, the strategy is to attempt to patent the system architecture in order to safeguard the platform against patent trolls. We will first provide a brief outline of the technology, and then dive into the strategic reasoning behind patenting this particular part of the project in terms the competitive and strategic advantages it would give the team. We will then talk about what the risks of not patenting are and how these hurdles can be overcome.

3.2.1 The Aspect: System Architecture

Being the platform that forms the intermediary between data collection and analysis mandates robustness, fault tolerance and ease of use. For widespread adoption, not only is it necessary to be first to market, it is also necessary to provide an easy to integrate solution. The application programming interfaces that would be provided should be seamless enough for people in the adjoining industries to build upon. The product's system architecture caters to all these requirements. We believe the innovation that went into meeting the constraints imposed by the requirements puts the team in good stead to warrant a patentable system architecture. This architecture allows for an instant state recall, clever storage and persistent reconnection as broad means to implement fault tolerance and robustness - aspects that are quite novel and developed from needs specific to handling sensitive health data.

3.3.1 Problems with Software Patents

Software patents are notoriously difficult to obtain due to the difficulty involved in proving the novelty and obviousness of them, unless they are extremely specific. Common problems include

1. Deeming software is math, and math is not patentable.(Jones, 2009)

2. Software programs being different than electromechanical devices because they are designed solely in terms of their function. While the inventor of a typical electromechanical device must design new physical features to qualify for a patent, a software developer need only design new functions to create a working embodiment of the program.(Plotkin, 2002)
3. Computers are deemed to be the "designers" and “builders” of the structure of executable software. Software developers do not “design” the executable software's physical structure, but merely provide the functional terms.(Plotkin, 2002)
4. The large number of micro-niches in software and relative few examiners mean patent examiners seldom have sufficient know how to recognize if the technology disclosed is innovative or novel enough.(Bessen, Meurer, 2008)

For the reasons stated above, it may not be possible for a small team of engineers to mount a serious patent application. In order to overcome this, we plan on leveraging U.C. Berkeley’s Office of transfer office, the Office of Technology Transfer, to help design and formulate a strong application to provide the best chance of being awarded a patent.

3.4.1 The Strategy

Applying for a patent is best in this case as it would protect the business from the threat of patent trolls and also would set a legal basis by which the business can protect its work. Forgoing patents could spawn copycat businesses that would patent this architecture, and could ultimately lead to the business being coerced into paying licensing fees for technology it originally developed. Applying for the patent protects the product irrespective of the outcome of the application for two reasons:

1. Approval of the patent will ensure that other competitors do not develop similar technologies, and thus raise barriers to entry. While it is true that the business may ultimately pursue an open-source strategy to encourage an organic growth of a strong third party community (which in itself raises barrier to entry), being awarded the patent ensures that the economic moat the team aims at building is deeper and wider.

2. Failure to get approval of the patent reassures the business that the same technology likely will not be awarded to a patent troll, giving peace of mind.

It is true that established competitors are much larger and financially capable than a startup this capstone group can produce, but the playing field can be leveled by leveraging the resources of the university to apply for a patent. This will both ease the financial burden, and cast a much larger enforcement net over the patent then could be done alone.

3.5.1 Risks of Not Patenting

If the backend architecture of the platform is not patented during the development phase, the team would leave itself open to the threat of having the architecture stolen and patented by another entity. As a consequence, the team could lose revenues from licensing, or even have to pay for technology the team invented in the first place. This could happen if some other group invented similar back-end architecture and patents it before the group's project is finished and released as an open source software. In addition, due to the highly collaborative nature of these projects, there is a risk that external developers would file a patent for the technology themselves. This would distract the team with crippling legal proceedings rather than focusing on further innovations.

Because of the risks of not patenting, it is advisable that patentable parts of the project are quickly pursued. To achieve this, an internal search utilizing university resources such as the Office of Technology Licensing (OTL) must be conducted to discover any prior art that may stand in the way of granting a patent. Legal experts should be then involved as quickly as possible to explore the potential patentability of the system architecture. Upon a thorough literature review by these experts, their advice will be taken and put into action.

In the case that it is found that a patent would be unlikely to be granted, the group would focus on making the platform open source as soon as possible in order to reduce the risk of infringement of patent rights. This establishes a record of prior art which can be used as a legal defense against patent trolls. A dual-licensing model will then be adopted to produce revenue. With dual-licensing, clients can use the technology under the condition that they redistribute their technologies as open source as well. If these entities wish to commercialize technologies utilizing this technology however, they will be required to pay a licensing fee on a per-use basis.

In conclusion, patenting the backend architecture provides a shield from legal problems and enables safe cooperation with external developers to speed up the development of the platform. In the unfortunate case that the patent is not granted, the only way to reduce legal risk is to release the platform as quickly as possible to establish prior art against potential litigation.

4.1.1 References

Centers for Medicare & Medicaid Services.
"National Medicare Readmission Findings: Recent Data and Trends." *Office of Information Product and Data Analytics* (2012).

Alnowami, M., et al.
A quantitative assessment of using the Kinect for Xbox360 for respiratory surface motion tracking. *SPIE Medical Imaging*. International Society for Optics and Photonics, 2012.

Banet, Matthew J., et al.
Body-worn sensor featuring a low-power processor and multi-sensor array for measuring blood pressure. U.S. Patent Application 13/346,408.

Banet, Matthew J., et al.
Vital sign monitor for cufflessly measuring blood pressure using a pulse transit time corrected for vascular index. U.S. Patent Application 14/072,305.

Bessen, James and Meurer, Michael J.
2008, "Patent Failure: How Judges, Bureaucrats, and Lawyers Put Innovators at Risk, Princeton University Press, http://press.princeton.edu/titles/8634.html, Accessed March 02, 2015.

Dolinay, Viliam, Lucie Pivnickova, and Vladimir Vasek
Objectivization of traditional otoneurological examinations based on Kinect sensor Hautant's test based on Kinect. *Control Conference (ICCC), 2014 15th International Carpathian*. IEEE, 2014.

Ghosh, Rajib; Heit, Juergen; Srinivasan, Soundararajan
Telehealth at scale: The need for interoperability and analytics
International Conference on Information and Knowledge Management, Proceedings, p 63-66, 2011, *CIKM 2011 Glasgow: MIXHS'11 - Proceedings of the 1st International Workshop on Managing Interoperability and Complexity in Health Systems*; Accessed February 16, 2015

Hall, Joseph L., and Deven McGraw.
For telehealth to succeed, privacy and security risks must be identified and addressed. *Health Affairs* 33.2 (2014): 216-221.

Hines, A. L., Barrett, M. L., Jiang, H. J., & Steiner, C. A.
Conditions with the largest number of adult hospital readmissions by payer, 2011 (2014). http://www.hcup-us.ahrq.gov/reports/statbriefs/sb172-Conditions-Readmissions-Payer.pdf,
accessed February 16, 2015

Jones, Pamela
2009, "An Explanation of Computation Theory for Lawyers"
http://www.groklaw.net/pdf2/ComputationalTheoryforLawyers.pdf Accessed March 02, 2015.

Lars, Nile
2014 Connected Medical Devices, Apps: Are They Leading the IoT Revolution — or Vice Versa?
Wired
http://www.wired.com/2014/06/connected-medical-devices-apps-leading-iot-revolution-vice-vers
Accessed Feb 27, 2015

Miller, Edward Alan.
"Solving the disjuncture between research and practice: telehealth trends in the 21st century." Health Policy 82.2 (2007): 133-141.

Morea, Stephan
2014 IBISWorld Industry Report OD5775: Telehealth Services in the US. http://clients1.ibisworld.com/reports/us/industry/default.aspx?entid=5775, Accessed February 16, 2015

Open Source Initiative
Welcome to The Open Source Initiative. http://opensource.org , accessed February 16, 2015

Plotkin, Robert
Intellectual Property and the Process of Invention: Why Software is Different, 2002, International Symposium on Technology and Society, http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=1013821, Accessed March 02, 2015

Porter E., Michael
2008 The Five Competitive Forces That Shape Strategy. Harvard Business School Publishing Corporation. https://hbr.org/2008/01/the-five-competitive-forces-that-shape-strategy. Accessed February 16, 2015

Skala, Karolj, et al.
4D thermal imaging system for medical applications. *Periodicum biologorum* 113.4 (2011): 407-416.

Swan, Melanie
2012 Sensor Mania! The Internet of Things, Wearable Computing, Objective Metrics, and the Quantified Self 2.0, Journal of Sensors and Actuators, Page 217-25, http://www.mdpi.com/2224-2708/1/3/217/htm#sthash.e0RU9ljv.dpuf, Accessed Feb 27,2015

Zhang, Zhengyou.
Microsoft kinect sensor and its effect. *MultiMedia, IEEE* 19, no. 2 (2012): 4-10.

(END OF COMMON REPORT)

Expanded Tele-Health Platform for Android

Final Capstone Report

Reviewers:

Dr. Ruzena Bajcsy

Dr. Ali Javey

Daniel Aranki

Jochem van Gaalen

25926165

May 12 2015

For Mom, Dad and my Brother and Sister

Executive Summary

Each year, $17 billion dollars is spent on patient readmissions. The most preventable readmission factors - and likely the most severe - are improper administration of therapy and medication. To combat this, we are developing a lean, smartphone-based telemedical system that not only provides personal health statistics to the patient, but relays this information to the doctor allowing for a more tailored approach to medicine. By leveraging the latest open source health monitoring algorithms and libraries, we have created a reliable, low-cost and easily implementable solution that has the potential to simultaneously improve patient care and dramatically reduce healthcare expenditure.

The team, advised by Dr. Ruzena Bajcsy and led by Ph.D. student Daniel Aranki, consisted of several members, including four Master of Engineering students (Phillip Azar, Adarsh Mani, Quan Peng, and Jochem van Gaalen) specializing in Robotics & Embedded Software concentration in the EECS department, and undergraduate engineering students Priyanka Nigam, Arjun Chopra, Sneha Sankavaram, Maya Reddy and Qiyin Wu.

Over the course of the academic year, an application in Android was built and tested through laboratory validation of individual components, and a pilot study conducted near the end of the project. The results of this work is outlined in this report, along with challenges faced and potential next steps for this project.

Table of Contents

Introduction

Hospital readmissions are a major economic burden for the healthcare system, with approximately $30 billion spent annually on chronic heart failure (CHF) patients alone[i]. Several studies have shown that tele-monitoring of these patients has various benefits, including reduction of readmission rates [ii] [iii] [iv] [v] [vi] [vii].

Our project aims to develop and test an open-source, highly configurable Android platform that allows doctors to monitor and determine risk of re-admittance for patients with chronic health conditions released from the hospital. This platform will utilize various sensors such as accelerometers, gyroscopes and cameras already onboard most smartphones to relay vital sign data (such as heart rate, blood pressure, temperature and caloric expenditure) to both patients, and to their assigned healthcare professionals in a privacy aware, secure and fault-free manner.

The group's specific contribution to this project was the development and testing of novel vital sign extraction algorithms for heart rate and blood pressure using the previously developed fault-tolerant framework. The group also expanded the communication functionality of the application by implementing a fault-tolerant Bluetooth framework to help interface with currently available medical devices such as blood pressure monitors.

Team Tasks

Because of the modular nature of the capstone project, the group was divided such that tasks could be completed in parallel. This ensured that the group stayed efficient throughout the semester, and could rapidly gain knowledge in the various aspects of the project. The task descriptions for each part of the project, and their main individual contributors are outlined below.

Bluetooth

The Bluetooth portion of the project involved creating an abstract architecture for Android such that Bluetooth enabled devices could easily interface with the overall architecture of the data-collection system on the phone. Fault-tolerance was also built into this interface. This portion of the project was primarily handled by Phillip Azar and Adarsh Mani.

Heart Rate

Vital sign extraction was another large portion of the group's project. Some work was previously done in developing algorithms for extracting caloric expenditure, a measure proposed to be a good indicator of a person's fitness level, but some other key measures, such as heart rate and blood pressure remained unimplemented. As such, the other half of the group, (Quan Peng and Jochem van

Gaalen) developed algorithms that could provide measurements of heart rate and blood pressure in real-time, with minimal extra hardware. This task was further divided such that Quan Peng worked on developing a heart rate algorithm based off variations in facial color from a camera video, and I worked on extracting blood pressure (BP) through a combination of photo plethysmography (PPG) and auditory sounds.

Integration of Parts

While the role of each segment of the project remained quite tangential to each of the other parts, these segments were recombined upon their completion to produce an Android app that could be used in pilot studies. For example, both heart rate algorithms would be used against one another to determine cases in which one would be preferable over another. Furthermore, the Bluetooth protocol would be utilized to allow seamless integration of other sensors such as blood pressure monitors into the study. This study, conducted at the conclusion of this capstone project, aimed to validate the individual components developed during this year, and also measured performance metrics such as power consumption of the app, speed, processing time across different phones, data transmission characteristics, and usage participation. These metrics were be logged with the hope of improving the app in future iterations and identifying next-steps for app extensions. Future studies will test fixed bugs, new features and overall usability.

Literature Review

At the beginning of the project, a state of the art survey was conducted of both academic research literature and commercially available products to determine an appropriate initial candidate algorithm for BP measurements. Focus has largely shifted from improving measurement accuracy of single-measurement modes to pursuing continuous measurement methods[viii]. These continuous measurement methods target the development of wearable health collection systems, which are largely thought to be the harbingers of the upcoming push towards wearable devices.

The following is a summary of the major measurement methods suitable for implementation in mobile applications:

1) Polyplethysmography + Machine Learning[ix]
   - A camera-based method that uses the discerns heart rate from the lightening and darkening of blood in the capillaries of 'thin' body parts such as fingers. Machine learning algorithms are then applied to incorporate other metrics such as weight and age into a reasonably accurate estimation of blood pressure.
2) Polyplethysmography + Stethoscope Audio Recording[x]
   - Utilizes the same camera method as described before, but combines audio data with this data to estimate blood pressure.
3) Sphygmomanometry[xi]
   - Most widely used method of measuring blood pressure. Utilizes a cuff to press a microphone against the arm, and listen for specific sounds made by the heart, commonly referred to as Korotkoff sounds.
4) Doppler Radar[xii]
   - Makes use of sonic waves pulsed through a vein and reflected back. Blood pressure can be determined by the transit time of these waves.
5) Electro-Pneumatic Control Loop[xiii] [xiv]
   - An impedance-based system comprising of several control loops that adjust for quick and long-term changes in blood pressure. This system is a more advanced version of (1).
6) Electromechanical Film[xv]
   - Utilizes a film that, when placed over the heart, can convert mechanical vibration of the heart into an electrical signal.

Upon completion of literature review to become familiarized with the state of the art, a review of commercially available products was conducted. The results of this search are summarized below. Once again, it should be noted that this list is by no means exhaustive, but rather highlights major BP collection methods currently commercially available.

1) Withings Blood Pressure Monitor (Smartphone Compatible, Cuff-Based)
2) Omron BP-652 Automatic Wrist 7 Series Blood Pressure Monitor (Standalone, Cuff-Based)
3) HealthStats Wireless Wrist Monitor (Force Transducer-Based)
4) Visi Mobile Finger Monitor (Optical-Based)

From this review, a BP measurement method was selected for an initial implementation. In order to narrow the search, the group decided that the first implementation of BP measurement should use minimal additional sensors to what is already available on the Android smartphone. Furthermore, preference was given to a measurement system that was relatively straight-forward to implement. The reasoning behind this being that as development progressed, increased familiarity with the Android programming environment (Eclipse) would allow for faster and more advanced modifications to the chosen algorithms such that it would function in real-time. As such, it was decided that the measurement method utilizing polyplethysmography and stethoscope audio recording presented by Chandrasekaran et al. would be implemented.

Summary of Chosen Algorithm

In the paper presented by Chandrasekaran et al, an offline algorithm was developed for determining blood pressure using a combination of Korotkoff sounds (heart sounds) recorded by a microphone, and a video recorded by a cellphone camera with its flash turned on. These measurements are collected on a cellphone and then processed using an offline, MATLAB algorithm on a computer. The paper shows that accuracies ranging from 94-99% can be achieved in measuring blood pressure can be achieved under laboratory conditions.

Methods and Materials

Approach

To develop an accurate, real-time estimator of heart rate, and subsequently, blood pressure, a careful milestone based method was adopted. Each phase of development required the addition of slightly more features starting with basic, offline algorithm validation until the full real-time algorithm was built including all of the safety checks. These phases also included gradual movement from the research and discovery centric environment provided by MATLAB, to the final implementation in Android. The end of each phase included a series of laboratory tests to confirm that the results gathered were still accurate and required minimal computing time.

The purpose of following this type of development was to quickly identify any potential problems with the validity of the algorithm. Also, due to unfamiliarity with programming in the Java-based Android environment, performing a transition from a known environment (MATLAB), to Android would allow for faster validation and production of initial results.

Instruments

Because of the nature of app development, relatively few instruments were needed when compared to other capstone projects. In total, a PC and an android based phone were used for the bulk of development, while an Omron blood pressure monitor was used for algorithm validation. Later on, stethoscopes, guitar pickups and external cellphone microphones were used to test different methods of capturing auditory signals of the heart.

From the start of the capstone projects, the aim was to keep the number of peripheral devices needed in addition to a phone to a minimum such that users of the tele-health monitoring system would not be inconvenienced by extra devices.

Experiments

Most experiments run in the algorithm validation phase of the project involved taking simultaneous measurements of heart rate using an Omron blood pressure monitor. These experiments involved taking multiple measurement of heartrate using different parameters (such as measure time, different measuring areas etc.) to evaluate accuracy, and the effect that changing these parameters had on the final measurement. The results of these experiments are explained in detail later on in this report.

Results and Discussion

Blood Pressure

My specific contribution to this project has been to implement a blood pressure measurement algorithm as presented by Chandrasekaran et al. using the previously developed Android-based fault-tolerant framework designed for vital sign monitoring on Android-equipped devices. This method of measuring blood pressure (BP) involves utilizing a combination PPG and Korotkoff sounds to estimate the user's blood pressure.

To do this, the user must place his or her finger over the cell phone camera while the flash LED is turned on and take a video of approximately 10 – 20 seconds in duration. By using the flash, the finger is illuminated. When the heart pumps, blood fills the capillaries of the finger and increase the finger's reflectance, making it appear brighter red. Conversely, after a heartbeat, blood leaves the capillaries and causes the finger to absorb more light (making it appear darker).

The major innovation is that the algorithm will be modified such that it works in real-time, and on mobile devices: an aspect of vital sign extraction that is underdeveloped in current literature. Furthermore, additional algorithms will be developed such that the data collected can be safe-guarded against invalid inputs or tampering. Again, the goal is to implement all of this in real-time, which is a challenge considering the variability between the processing and hardware capabilities of the plethora of Android phones currently on the market. The goal is to be able to provide doctors with a direct comparison of different data collection methods associated with BP measurements, and then also to use the collected data as a component of future patient readmission risk assessment analysis.

The following section outlines the major milestones reached throughout the development of the heart rate extraction algorithm and development of the Android app. The progress made at each major revision of the algorithm will be outlined in detail, and also the challenges faced at each of these steps is outlined.

MATLAB

The first step in measuring BP described in the Chandrasekaran paper is to measure heart rate (HR). To validate this step, a HTC One X smartphone was used to take various videos of the index finger at different resting heart rates. The videos were loaded in MATLAB, and the average red pixel intensity of each frame was computed sequentially over entirety of the video to give an averaged time-series of the video. Figure 1 below shows a plot of the averaged red pixel time-series data across a sample video. Note that the first few seconds of the video data has been removed as a preliminary data cleaning step to remove the undesired effects of camera focusing.

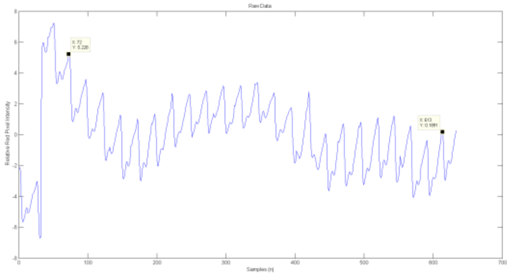

Figure 1: Mean Red Pixel Intensity

From Figure 1, it is clear that heart rate data can be extracted from this data. Upon visual inspection, there exist 22 peaks over 18.5987 seconds of video. This translates to a HR of about 70.97 beats per minute (BPM) which matches the self-measured value range of 70 BPM (± 4 BPM).

The next step of the data cleaning process is de-meaning and band-pass filtering the data to remove unwanted noise and isolate the HR frequency. The band pass filter used was a Butterworth filter with stopbands of 30 and 140 BPM and a passband from 50 to 120 BPM. Figures 2 and 3 show the results of performing both these operations on the data. Figure 2 shows the de-meaned mean red pixel intensity data in red, and the filtered data in blue. Figure 3 shows the spectral frequency composition of the filtered data using Fast Fourier Transform.

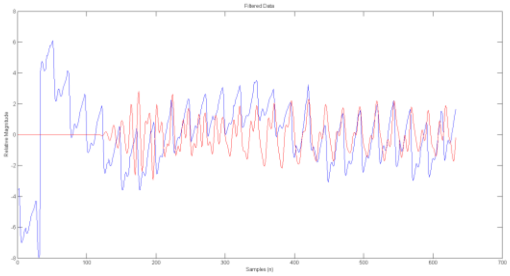

Figure 2: De-Meaned Data (Blue) and Filtered Data (Red)

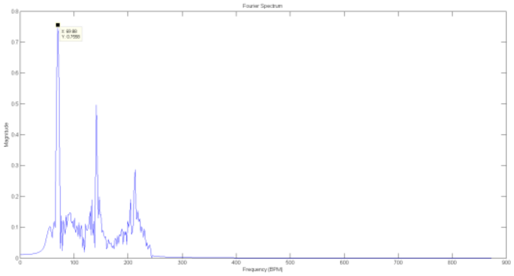

Figure 3: Spectral Composition of Filtered Data

We can see from Figure 3 that the frequency with the greatest magnitude is 69.88 BPM which matches the self-measured value of 70 BPM (± 4 BPM).

Challenges Faced

Upon completion of performing frequency analysis on a sample video, several additional videos were captured and analyzed in the same manner as previously described. One video in particular showed that the frequency makeup of a signal can have two peaks of close proximity which can cause confusion about which frequency is actually the true signal. Figures 4 and 5 show the result of one particular case in which it is not immediately clear which frequency corresponds to the real HR.

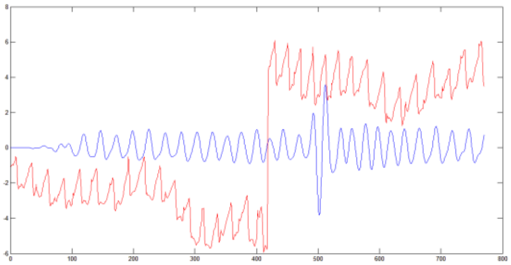

Figure 4: De-Meaned data (red) and filtered data (blue)

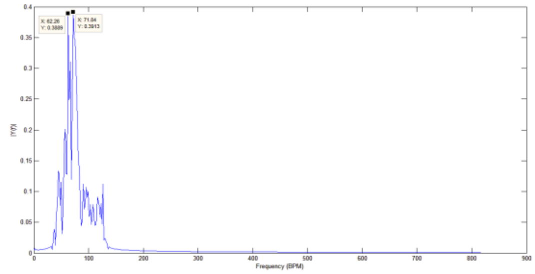

Figure 5: Spectral composition of the filtered data in Figure 4

As we can see, this introduces uncertainty into the measurement. In order to resolve this issue, a 5-point moving average filter was implemented in the next step of development, done in Java.

Java Development

In this phase, a few minor tweaks were made such as relaxing the bounds on the band-pass filter (42 Hz – 240 Hz pass-band) used such that it could accommodate higher heart rates without cutting them off. Because this change introduced sporadic instances of higher frequencies in the signal, a 5 point moving average filter was then be applied to the signal to reduce the effect of these higher frequencies.

In addition, an attempt was made to implement the heart beat interval extraction algorithm developed by Pan and Tompkins, but it was found that due to the low sampling rate of the camera, this method did not provide usable results. As such, the data was simply converted to the frequency domain via FFT and the predominant frequency contained in the data was used as the heart rate. As is discussed later, this was found to be an adequate method of extracting heart rate information.

Laboratory Tests

In the last phase, the MATLAB portion of the PPG extraction was completed. After this was done, a series of lab tests were run comparing the results of this algorithm to an accepted heart rate measurement device. In this case, the "ground truth" data was supplied by an Omron blood pressure monitor with the capability of supplying heart rate information. A total of 15 tests were conducted where both PPG and ground truth data were collected simultaneously. The video data was then used to compare the developed algorithm in MATLAB to the ground truth. Please note that in all cases, the algorithm developed discarded the first 110 frames of the video, as each video proved to have noisy information for roughly the first 110 frames due to camera focusing. It was found from these tests that the results largely agreed with the ground truth data. The exact data is shown and discussed in detail later.

Using the apache commons library for Fast Fourier Transforms and JFreeChart for plotting (which is only used for checking data correctness) the full algorithm was implemented in Java. An example plot of the algorithm output is shown below.

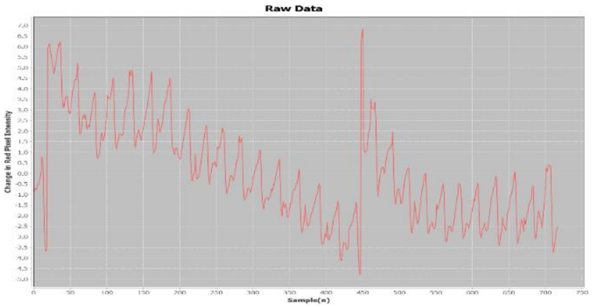

Figure 6: Raw Data plot in Java

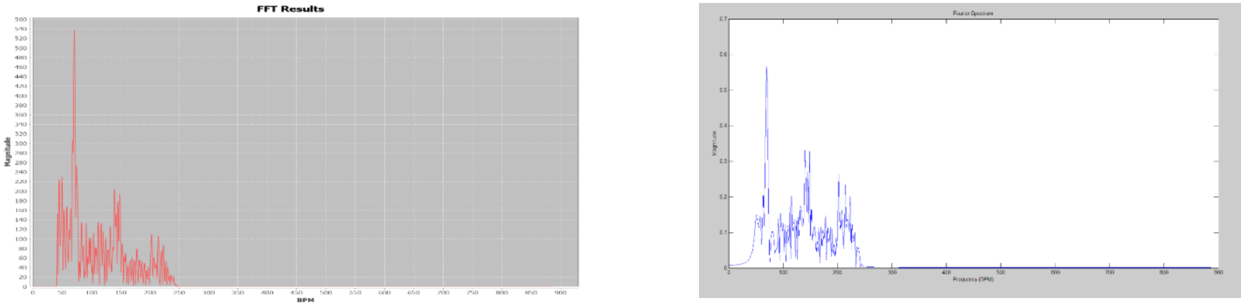

Figure 7: FFT Results in Java (Left), and in MATLAB (right)

The results from plotting appeared nearly identical to those generated in MATLAB. After this was complete, the algorithm in Java was compared to the MATLAB and ground truth results. Below is a summary of the results.

Table 1: Comparison of Datasets

Test No.	**Omron**	**FFT Java**	**FFT MATLAB**
1	67	66.69	66.69
2	73	70.97	69.24
3	67	63.64	63.86
4	74	74.08	72.36
5	71	70.52	69.88
6	75	70.23	70.02
7	70	69.98	57.85
8	75	69.51	69.71
9	74	72.21	73.27
10	73	71.65	69.94
11	70	69.31	68.45
12	72	71.39	69.2
13	67	73.89	72.39
14	68	69.72	70.58
15	68	68.81	70.08

From here we can see that both algorithms largely agree with one another, but there appear to be small differences in the maximum reported frequencies. This occurred because a static filter was implemented in Java while MATLAB recomputed its filter each time based on the reported sampling frequency of the video.

One test (no. #7) produced a large difference between the reported MATLAB results, and those produced in Java. The output plots for both are shown in the figures below.

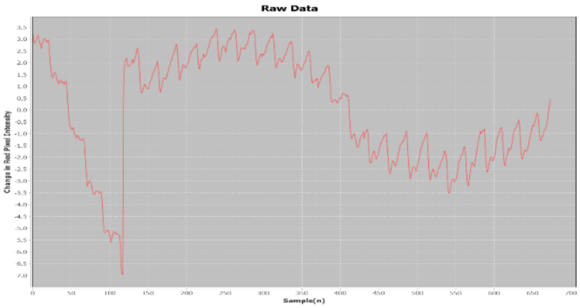

Figure 8: Raw Data Plot in Java

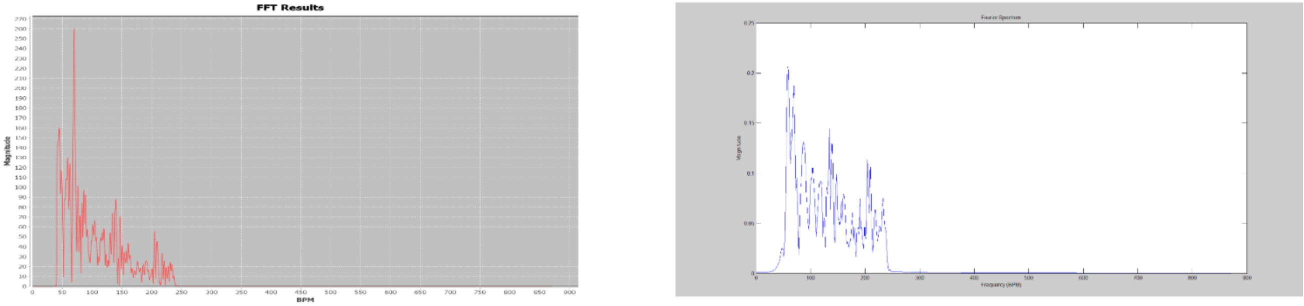

Figure 9: FFT Results in Java (Left) and in MATLAB (Right)

At the time of writing, it remains unclear what the underlying cause is behind the large difference in reported maximum heart rate or why the frequency spectrums differ greatly between the two plots. The most likely reason is because there was a bug in the original implementation of the convolve algorithm in which the signal to be convolved was not reversed prior to multiplying the two signals. Fixing this issue brought correct results.

Data Analysis

From the data presented, the correlation coefficients and RMS errors of all three algorithms developed were calculated. Note that the 'peak' algorithm simply calculated local maxima across the dataset and was used as a first attempt to extract heart rate information before the frequency analysis approach was attempted. The results are shown in Table 2.

Table 2: Statistical Summary of Developed Algorithms

	Java	**MATLAB**	**Peak**
Correlation	0.480	0.324	-0.298
RMS	2.615	4.248	11.042

From these results, it is clear that the correlation coefficients across all three algorithms show low correlation with the ground truth. The reason for this is because the heart rate was not varied when the tests were performed. As such, only the correlation between the noise of the ground truth and the algorithms are found by calculating correlation.

On the other hand, it is clear that the RMS error has improved drastically across each iteration of the algorithms developed. The average RMS error using the algorithm developed in Java is now just 2.804 bpm.

Code Design Changes

After this was complete, the code written in Java was then divided up into individual methods. The functionality of these methods was increased such that future work can easily implement past methods. For example, the n-point averaging method was changed such that a user could get averages across different steps sizes.

Another addition made was adding exception handling. This was done such that the code would be more robust and less prone to unexpected errors.

Finally, the code written to extract heart rate was packaged into a single class which accepts an array representing average red pixel intensity and a sampling frequency and returns a single heart rate value. This was done such that the heart rate extraction algorithm could be used in a sample application for the M.Eng capstone expo.

Algorithm Summary

With implementation of the heart rate extraction algorithm complete from the point at which red pixel data has been extracted, we can now provide a summary of the algorithm developed so far. Figure 10 summarizes the steps taken in order to calculate heart rate, and save this data for storage and data transmission.

Figure 10: Heart Rate Extraction Algorithm Summary

Note that when a data point is ready to be saved, it is stored in an encapsulator which is a thread-safe datatype which stores a copy of the data added to it, along with a timestamp of when the data was added.

Real-Time Development

Now that the code was ported to Java, work was then focused on real-time implementation of the algorithm.

Lazy Video

The first attempt made at real-time processing of video was to implement was is known as a "lazy video" class. This class stored the video recorded in memory without additional processing. Once the video was finished being recorded, the frames would be processed one at a time, and the heart rate would be calculated in the manner previously discussed. The problem encountered with this type of system architecture is that loading each frame into memory is time consuming, and is by far the largest bottleneck of the system. From a user's perspective this is akin to taking a video to be analyzed, and then having to wait many times more than the actual length of the video to receive a final measurement from the algorithm. This is very user-unfriendly especially if invalid data is collected and the user must take the measurement again. Furthermore, the system is limited by the fact that Android does not allow direct access to individual video frames, and instead offers a function that will return a frame that is closest to a specified time. This introduces undesirable uncertainty and error into the estimates taken, but does not totally invalidate the potential of the system. To solve the aforementioned fetching issue, an attempt was made to cache the frames for faster fetching. Upon testing however, this was found to offer no substantial increase in processing time that would allow for fast processing. Finally, storing the video locally on the phone was also memory intensive, which is not desirable from an app usability standpoint. Thus, a new processing architecture was required.

On-the-Fly Processing

The method that was adopted for real-time processing involved simply processing each frame of the camera on the fly. This was done by extracting the current preview frame stored on the camera. In essence, a preview frame is a dynamic snapshot of what the camera is currently looking at. Each time a new preview frame is generated, the average red pixel intensity would be extracted, and stored as a single number for processing after sufficient data points were collected. Note that each frame is discarded after the average red pixel intensity is calculated. After sufficient time has elapsed, the aforementioned algorithm is run on the entire dataset to produce a single number for heart rate. Once complete, more data points are added in the same manner as described, but this time older data is removed to keep the size of the dataset fixed. The result of this is a moving window over which heart rate measurements are taken.

One caveat encountered when implementing on-the-fly processing is that due to the scheduling of tasks on the phone, the data of the preview frames is not collected at regular intervals. Because the underlying assumption of the original algorithm is that the data is sampled regularly, the data collected could not be used as-is. Instead, upon processing the data, neighbouring data points were interpolated using spline interpolation and then resampled at regular intervals of 30 frames per second (FPS). With this issue resolved, the heart rate algorithm could now be run without issue.

Note that the resampling rate was chosen to be 30 fps to be well above the minimum frequency required to measure the fastest expected heart rates of about 240 BPM, which is near the upper limit at the rate of which an average healthy human heart can pump[xvi]. The Nyquist sampling theorem gives that, at 200 BPM, the minimum sampling frequency is twice the maximum expected signal frequency. In this case, we require a minimum sampling frequency of about 8 Hz.

The next step in development was to the optimal time at which to stop taking measurements, and to run the algorithm to determine heart rate. The main trade-off here is that a longer test can provide more accurate results, but in return, the processing time slows down considerably after lengthy measurement times. It was decided that at maximum, a measurement would not take longer in duration than a standard blood pressure measurement. From the equipment available in the laboratory, it was found that this amounts to about 25 seconds on average.

As a test to measure this real-time algorithm, the blood pressure monitor was used as the 'gold-standard' and compared directly against the phone. A series of 20 samples were conducted this time (using the same subject for each test), with two different tests run at different measurement times. The 'Android Short' test allowed the algorithm to collect data for 10 seconds, while the 'Android Ideal' test allowed for 20 seconds of data collection. The results of these two tests are shown in Table 2. Note that the previous test results are also included.

Table 3: Summary of Laboratory Test Results

Test No.	**Android Short RMS**	**Android Ideal RMS**	**Java RMS**	**MATLAB RMS**	**Peak RMS**
1	14.8225	0.5776	0.0961	0.0961	1.4641
2	16.0801	6.8121	4.1209	14.1376	67.8976
3	10.8241	6.2001	11.2896	9.8596	3.2761
4	17.7241	3.61	0.0064	2.6896	44.3556
5	20.6116	5.1529	0.2304	1.2544	1.4161
6	0.5929	6.9169	22.7529	24.8004	83.3569
7	1.8225	0.1521	0.0004	147.6225	147.6225
8	0.0841	0.0025	30.1401	27.9841	226.2016
9	1.0201	1.7161	3.2041	0.5329	34.9281
10	2.9241	0.6889	1.8225	9.3636	183.8736
11	19.8916	61.4656	0.4761	2.4025	65.1249
12	29.3764	1.2544	0.3721	7.84	1.3924
13	13.6161	0.0001	47.4721	29.0521	614.5441
14	0.0009	0.0841	2.9584	6.6564	57.3049
15	3.3489	0.0361	0.6561	4.3264	417.7936
16	0.0729	40.3225			
17	0.0025	7.7841			
18	0.0289	1.0609			
19	9.7344	0.5329			
20	2.7889	0.6889			
RMS	**2.88**	**2.69**	**2.80**	**4.25**	**11.04**

We can see here that while providing more time for measurement produces more accurate results, the gain is limited. The RMS, or root mean squared, error between the actual reported heart rate shows that real-time algorithm is accurate in both tests – only off by 2.88 and 2.69 BPM RMS. This test provided the final validation that this algorithm could work in real-time and produce highly accurate results. However, in order to be fully integrated in a study, safeguards needed to be developed such that measurements only start when a finger is detected to be in front of the camera.

Input Validation

To increase the usability of the algorithm, the camera preview must be validated such that the algorithm only starts collecting data once a finger has been detected to be in front of the camera. In order to validate the measurements taken, a novel approach was developed in which the entropy, or 'business' of the incoming image was calculated for two different cases. Put simply, 'entropy' is a measure of both the number and the frequency of the colors present. For example, if an image has only one color present, its entropy is zero. An image which is highly varied in both the number of colors present and their occurrences will have high entropy.

This property was used in the following manner: first, the entropies of the red pixels and a greyscale[1] version of the image were calculated. These values were then compared to one another. It was found that valid images (those which the user covers the camera with a finger) exhibit high ratios (>9%) of red entropy to greyscale entropy. This is because the red pixels are predominant in the image, and relatively low amounts of green or blue pixels are present to change the total number of unique colors present in the greyscale image. In a random invalid image, green and blue pixels have a much larger relative contribution (i.e. they cause a much larger variation of different grey pixel intensities present) to the greyscale image, and thus cause the entropy of the image to be larger. With this, the heart rate collection algorithm was complete. The figures below show the algorithm at work in a sample case.

Place Index Finger over camera

Finger Detected. Calculating...

HR: 82.8369140625 with 167 data points (in 7 seconds)!

No finger detected (Random Image) *Finger Detected, Gathering data* *Displaying Data & Gathering New Data*

Figure 11: Example Camera Preview Images for Validation Stages

---

[1] Note that greyscale contains information about all three colors of a pixel in one value. This was done by taking the average intensity of the red, green and blue values.

System Summary

With all of the individual components for real-time measurement of heart rate completed, we can now take some time to summarize the workings of the system as a whole and show how the different classes implemented within the system interact with one another.

At the start of measuring heart rate, a menu is displayed to the user showing the different options by which one can measure heart rate. This display consists of two different fragment user interfaces (UIs), one of which contains the list of options (on the left) along with the 'measure' button, and another which displays instructions on how to use each algorithm. The UI used in the pilot study is shown in figure 16.

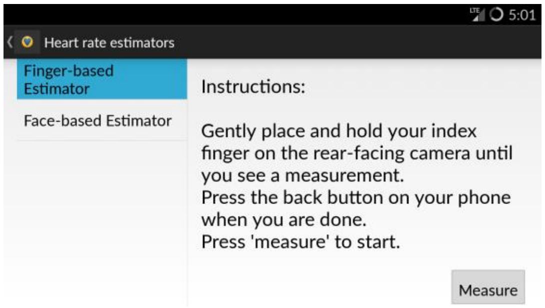

Instructions:

Gently place and hold your index finger on the rear-facing camera until you see a measurement.
Press the back button on your phone when you are done.
Press 'measure' to start.

Figure 12: UI Fragment Example

Upon pressing 'measure' the program then functions in the same way as has been outlined before. An encapsulator for storing the measured heart rate data along with the time stamps associated with that data is created, and then passed to a video estimator class which contains the main algorithm for calculating heart rate. This class instantiates a new fragment UI class which inflates to fill the screen and show the user the current camera preview. Upon completion of this, the fragment starts the processor, which acts as the 'hub' for requesting and processing data. The processor contains the finger detection algorithm, and also averages out the red pixel data when a finger has been detected and saves this and a time stamp to an array. Once sufficient time has passed, the processor interpolates the data, and then sends this back to the estimator which returns a single heart rate value along with a timestamp. This data point is then added to the encapsulator, and the entire process of calculating heart rate is then repeated until the user presses the 'back' button on their phone or takes their finger off of the camera lens. Lastly, while data is being added to the encapsulator, a listener registered to that encapsulator waits for new data to be added and promptly sends this over to a remote server. The figure below provides an outline of this system.

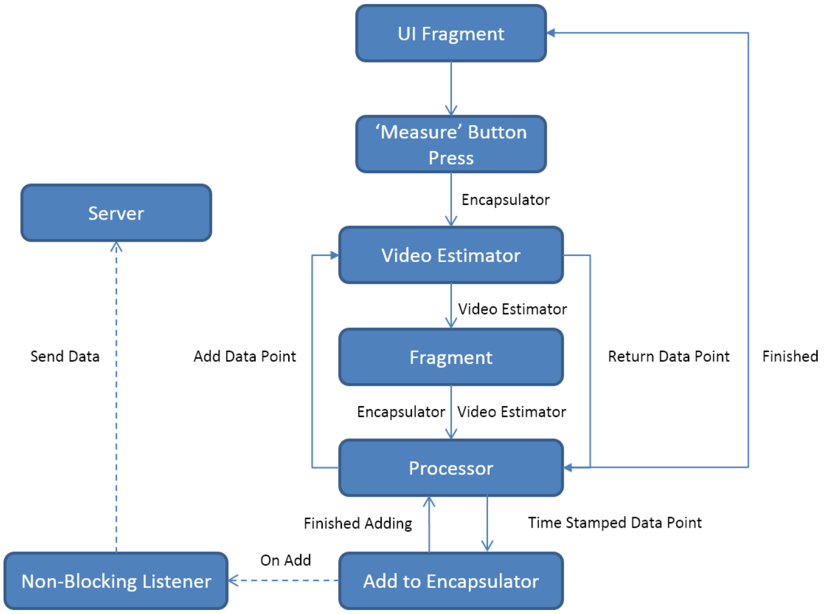

Figure 13: System Diagram

With this system successfully implemented in Android, the next step in development was to validate the algorithms developed amongst varying Android supported phones and users. To do this, a pilot study was conducted, which is discussed in the next section.

Pilot Study

In order to validate the system as a whole, the individual parts of the project were synthesized into one single application, and a small pilot study consisting of 11 participants was run. The pilot study consisted of providing the participants a copy of the full application on their phones (note: since the application was written for Android, participants were limited on the basis of owning an Android phone), and allowing them to use the app freely. The participants were also asked to come into the laboratory once during the study period to take measurements with a Bluetooth blood pressure monitor by A&D (UA-651BLE) and with the finger and face heart rate measurement algorithms. This was done to validate each algorithm against a ground truth device across different participants. It is important to note that for the measurements taken in the laboratory, all participants used the same phone. This was done because some bugs in the application prevented reliable data collection across the slew of phones used in the study. A questionnaire was also given at both the start and completion of the survey to evaluate the usability of the different subcomponents of the app, along with future prospects of tele-monitoring systems in general.

Data Analysis

From the data collected from the study, we can now validate the algorithms effectiveness across different test subjects. Below is a summary of the root mean square (RMS) error across all of the measurements (N=243) that were taken with an accompanying ground truth test. Below is an error histogram and table summarizing the results of the study. Please note that the 'raw' dataset includes all of the data collected, while the 'outlier removal' dataset has removed poor data points in which a noisy or invalid signal was captured. This is known because the resulting reported HR is given as the extreme lower bound of the band-pass filter (40-50 BPM).

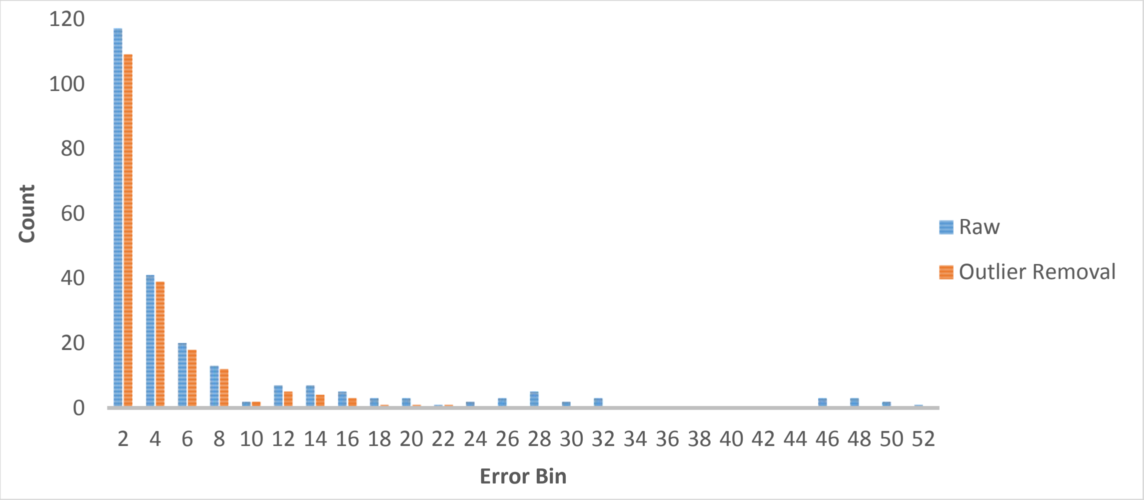

Figure 14: Pilot Study Error Histogram

Table 4: Pilot Study RMS

	Raw	**Outlier Removal**
Sample Count	243	195
RMS	13.52	4.89

As can be seen, the raw RMS is much higher than the previously tested RMS of 2.615 in a very controlled environment. From examining the raw data however, there appears to be a clear subset that is directly the result of a poor input signal. This is because the reported heart rate is at or near the cut-off frequency for the band-pass filter applied to the data. Removal of this outlying data reveals a significant improvement of the RMS error. Despite this however, the RMS error is still nearly twice that of the previously reported RMS. Further investigation of why this is so will be required in future pilot study, but there are a few factors that seem most likely, which are listed below.

1) Users are pressing their fingers too hard against the camera lens, resulting in degradation of the signal
2) Users are moving their fingers across the lens during data collection
3) Connectivity issues experienced throughout the study negatively affected the sampling rate of the data
4) The flash on the phone used for testing produces an image that is too bright/dim

These candidate problems will be tested more rigorously in future tests to determine the source of the additional error.

Areas of Improvement

From the analysis presented, there are a few areas in which the algorithm could be improved. First, in future pilot studies, more information should be collected so as to help pinpoint the cause of the erroneous measurements taken. Data such as the number of frames processed within the 20 second measurement window would help determine whether the algorithm suffered from a performance issue, or if the user was unclear as to how use the algorithm properly. Additionally, some extra safeguards can be put in place whereby low measurements (40 – 50 BPM) could prompt the user to redo the test. This would likely have an improvement on the quality of data received, but would make detecting actual heart rates between 40-50 BPM impossible. Finally, data that is validated using a Bluetooth blood pressure monitor should be packaged together for easier data analysis. Often, due to connectivity issues inside the room in which the tests were taken, results from the blood pressure monitor would be stored on the phone, and sent all at once after the phone established wireless connectivity some time later. Since the blood pressure monitor measurements were not time-stamped locally on the phone, the results of the finger heart rate tests (which were time-stamped) were at times difficult to sort between the tests used to validate the face-based heart rate algorithm and the finger-based one. In conclusion, future studies in the near term will modify this aspect of the study to allow for participants to use their own phones.

Qualitative Analysis

Along with the empirical measurement data, additional data was collected in the pre/post questionnaire about the user's views on topics such as their satisfaction with current health-monitoring devices, security, and overall interest in health-monitoring. Below are the results of the answer to the survey regarding measuring heart rate. The figures on the left show the pre-study questionnaire results while the figures on the right show the post-study questionnaire results.

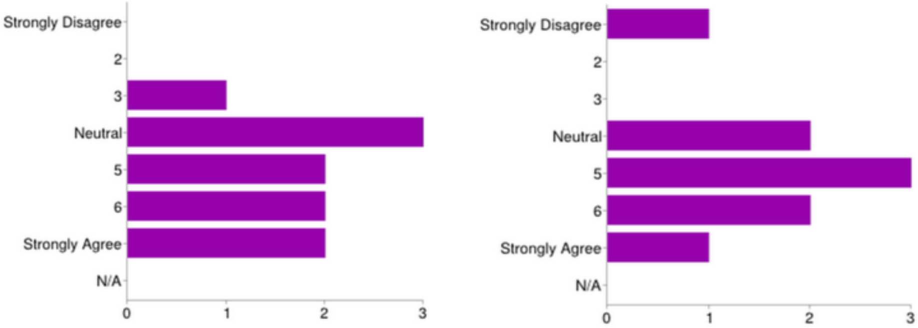

Figure 15: Question: I would measure my heart rate regularly if it were easy to do

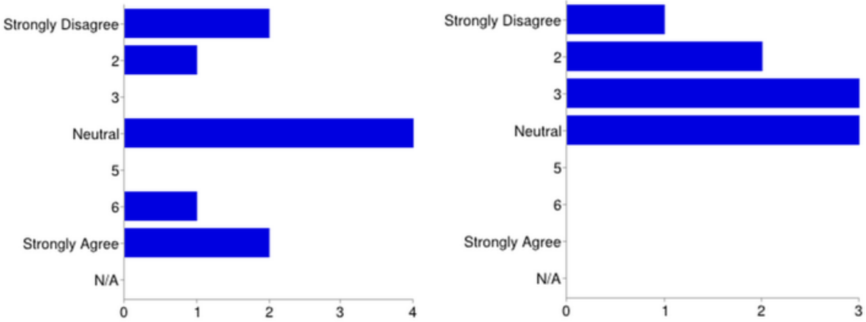

Figure 16: Question: I think non-cuff based heart rate measurements are inaccurate

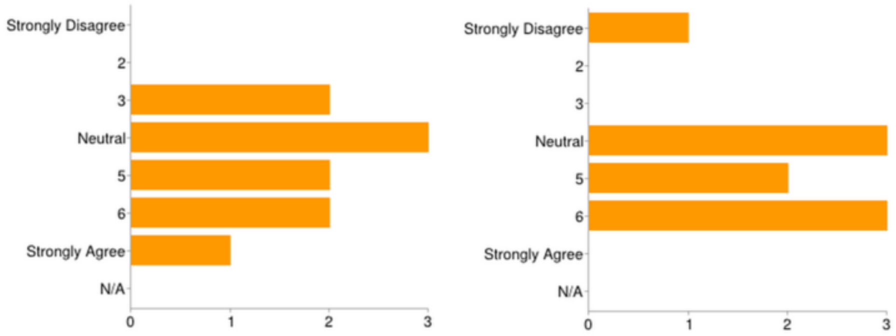

Figure 17: Question: I will use heart rate measurements to monitor my fitness

While the sample size of this survey is not sufficient to draw hard conclusions, we can look at these results qualitatively and use the results as motivation for more rigorous tests in the future. Overall, it seems that the pilot study was effective in changing the opinion of the subjects in a positive manner. After using the application, the subjects seemed more likely to measure their heart rate in the future, indicating that the application was relatively easy to use. Moreover, the subjects seemed to be convinced after the study that it is possible to accurately measure heart rate with non-cuff based technologies, after being split on the question beforehand. Finally, there seemed to be a slight decrease in favourability in the last question asked in the questionnaire. The cause of why this may be is unknown and should be investigated further. One hypothesis is that users saw their heart rate vary throughout the day and, while believing that the result was accurate, determined that measuring their heart rate may not be a good measurement to use when monitoring fitness levels.

These results come as somewhat of a surprise, as there were reports from the test subjects during the study that the application would crash at various points of usage. This included crashing upon opening the finger heart rate algorithm. It seems that despite this, the subjects saw the potential in the measurement algorithm and thus improved their answers in the second round of the survey.

Conclusion

This paper has presented a summary of the development of a real-time heart rate measurement on mobile devices using a vision-based method called photo plethysmography (PPG). Preliminary results have validated the accuracy of the algorithm, but additional work is needed to improve the performance of the algorithm across different devices. Some unrelated issues in the application that was built for the pilot study are hypothesized to be contributors to some of the data inaccuracy in the pilot study, but further detailed tests will be required to pinpoint the exact cause of this.

Upon the complete validation of this algorithm across multiple Android phones, work will shift to expanding the system to include real-time blood pressure measurements, which was the ultimate goal of this project from the onset.

---

[i] J. Xuan, P. Duong, P. Russo, M. Lacey, and B. Wong. The economic burden of congestive heart failure in a managed care population. *The American Journal of Managed Care*, 6(6):693, 2000.

[ii] S. I. Chaudhry, J. A. Mattera, J. P. Curtis, J. A. Spertus, J. Herrin, Z. Lin, C. O. Phillips, B. V. Hodshon, L. S. Cooper, and H. M. Krumholz. Telemonitoring in patients with heart failure. New England Journal of Medicine, 363(24):2301-2309, 2010.

[iii] R. A. Clark, S. C. Inglis, F. A. McAlister, J. G. Cleland, and S. Stewart. Telemonitoring or structured telephone support programmes for patients with chronic heart failure: systematic review and meta-analysis. *BMJ*, 334(7600):942, 2007.

[iv] M. Clarke, A. Shah, and U. Sharma. Systematic review of studies on telemonitoring of patients with congestive heart failure: a meta-analysis. *Journal of Telemedicine and Telecare*, 17(1):7-14, 2011.

[v] G. Giamouzis, D. Mastrogiannis, K. Koutrakis, G. Karayannis, C. Parisis, C. Rountas, E. Adreanides, G. E. Dafoulas, P. C. Stafylas, J. Skoularigis, et al. Telemonitoring in chronic heart failure: a systematic review. *Cardiology Research and Practice*, 2012, 2012.

[vi] S. Inglis. Structured telephone support or telemonitoring programmes for patients with chronic heart failure. *Journal of Evidence-Based Medicine*, 3(4):228{228, 2010.

[vii] C. Palaniswamy, A. Mishkin, W. S. Aronow, A. Kalra, and W. H. Frishman. Remote patient monitoring in chronic heart failure. *Cardiology in Review*, 21(3):141-150, 2013.

[viii] Pantelopoulos, A., and Bourbakis, N. A survey on wearable sensor-based systems for health monitoring and prognosis. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on* 40, no. 1: 1-12, 2010.

[ix] Visvanathan, A., Aniruddha S., and Arpan P. Estimation of blood pressure levels from reflective Photoplethysmograph using smart phones. *IEEE 13th International Conference on Bioinformatics and Bioengineering (BIBE)*, pp. 1-5, 2013.

[x] Chandrasekaran, V. et al. Cuffless Differential Blood Pressure Estimation Using Smart Phones. *IEEE Transactions on Biomedical Engineering*, vol.60, no.4, pp.1080,1089, 2013.

[xi] Perloff, D. et al. Human blood pressure determination by sphygmomanometry. *Circulation* 88, no. 5: 2460-2470, 1993.

[xii] Singh, A., Lubecke, V., and Lubecke, O. Pulse pressure monitoring through non-contact cardiac motion detection using 2.45 GHz microwave Doppler radar. A*nnual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC,* pp. 4336-4339, 2011.

[xiii] Penaz, J. Photoelectric measurement of blood pressure, volume and flow in the finger. *Digest of the 10th international conference on medical and biological engineering*, vol. 104, 1973.

[xiv] Fortin, J., et al. Continuous non-invasive blood pressure monitoring using concentrically interlocking control loops. *Computers in biology and medicine* 36, no. 9: 941-957, 2006.

[xv] Noh, S. et al. Ferroelectret film-based patch-type sensor for continuous blood pressure monitoring. *Electronics Letters*, 50(3):143, 2014

[xvi] Strobel, Ronald E., John D. Fisher, Gillian Katz, Soo G. Kim, and Anthony D. Mercando. "Time dependence of ventricular refractory periods: Implications for electrophysiologic protocols." *Journal of the American College of Cardiology* 15, no. 2 (1990): 402-411.