

RESEARCH DIRECTIONS IN CLOUD-BASED DECISION SUPPORT SYSTEMS FOR HEALTH MONITORING USING INTERNET-OF-THINGS DRIVEN DATA ACQUISITION

Alex Page, Shurouq Hijazi
University of Rochester
{apage4,shijazi}@ece.rochester.edu

Dogan Askan, Burak Kantarci
Clarkson University
{askand,bkantarc}@clarkson.edu

Tolga Soyata
SUNY Albany
tsoyata@albany.edu

Abstract

The Digital Health (D-Health) era is expected to be the “next big thing” since the invention of the internet, characterized by inexpensive and widespread medical data acquisition devices, widespread availability of identity-removed health data, and analytics algorithms that provide remote health monitoring feedback to doctors in realtime. Recent years have brought incremental developments in three key technological areas towards the realization of the D-Health era: data acquisition, secure data transmission/storage, and data analytics. i) For data acquisition, the emerging Internet-of-Things (IoT) devices are becoming a viable technology to enable the acquisition of remote health monitoring data. ii) For data storage, emerging system-level and cryptographic mechanisms provide secure and privacy-preserving transmission, storage, and sharing of the acquired data. iii) For data analytics, emerging decision support algorithms provide a mechanism for healthcare professionals to base their clinical diagnoses partially on machine-suggested statistical inferences that rely on a wide corpus of accumulated data. The D-Health era will create new business opportunities in all of these areas. In this paper, we propose a generalized structure for a D-Health system that is capable of remote health monitoring and decision support. We formulate our proposed structure around potential business opportunities and conduct technical feasibility studies.

Keywords: remote health monitoring; medical decision support; Internet of Things (IoT); visualization; analytics;

1. INTRODUCTION

The unprecedented growth in the Internet of Things (IoT) technologies makes it possible to talk about 50 billion connected devices through the internet by 2020 (Fernandez & Pallis, 2014). Among these devices are body-worn sensors that monitor personal health conditions. There has been a growing interest in wearable sensors in recent years and an emerging set of new products are commercially available (Jawbone, 2016; FitBit Inc., 2016; Apple Inc., 2016) for activity recognition, personal health monitoring, and fitness. For clinical use, long-term patient monitoring and management has also been considered by researchers (Pantelopoulos & Bourbakis, 2010; Son & et al., 2014; Page, Kocabas, Soyata, Aktas, & Couderc, 2014; Paradiso, Loriga, & Taccini, 2005; Milenkovi, Otto, & Jovanov, 2006; Istepanian, Sungeor, Faisal, & Philip, 2011; Soyata T., 2015). IoT-based data collection and cloud-based analytics are the driving factors of this technology as detailed in (Hassanaliyagh, et al., 2015). A doctor can prescribe a 2–3 day period of continuous physiological monitoring of a patient using low-cost wearable devices before a patient’s periodic physical examination. This monitoring data can be transmitted to the database, linked with the health records of the patient. Statistical inference algorithms can compare this patient’s data to a large database of other patients and provide the doctor with a rich set of suggestions. These

machine-inferred suggestions are invaluable tools which use technology for the benefit of human health.

The Digital Health (D-Health) vision described in the preceding paragraph promises to be a disruptive technology for human healthcare. In addition to saving the hospitals money, this type of decision support could improve diagnostic accuracy and might create third party business opportunities. However, before this vision can be fully realized, a set of challenges that need to be addressed are: (i) The privacy and security of the acquired data need to be ensured during its acquisition, storage, and processing. (ii) A large dataset for specific health conditions takes time to build and the accuracy of many decision support algorithms depend on the size of the database, thereby creating a natural vicious cycle. (iii) Despite being full aware of its potential, hospitals will be slow in embracing the D-Health concept due to the risks implied in basing decisions that can effect human lives on machine suggestions. (iv) It is not clear how this technology can turn into business opportunities. (v) The IoT technology is still in its infancy and it is not clear whether this technology will enable a secure and reliable sensing platform. (vi) Even if the data can be acquired reliably, it is not certain whether this data can be visualized in a non-overwhelming summarized format to be useful to the doctors and be embraced by them. (vii) Since large databases for many diseases are proprietary or simply do not exist, it is not clear whether statistical




Figure 1. Layers of the proposed remote patient monitoring system that is based on an IoT-Cloud architecture. Based on the challenges described in Section 1, as well as the available business opportunities that will be described in Section 3, it suffices to conceptualize the system as two super layers: The Front End represents the hardware and software, necessary for the secure acquisition of the patient health data. The Back End represents the cloud infrastructure to store and process the data, as well as the visualization and analytics algorithms running in the cloud.

inference is possible for a wide variety of diseases that can be detected through remote health monitoring.

In the rest of this paper, we aim at providing answers to challenges (iv)–(vii). Towards that end, we introduce a generalized system structure for remote health monitoring based on recent research directions, as well as our predictions in Section 2. In Section 3, we address challenge (iv) and identify a clear list of existing business opportunities. In Section 4, we identify the technical components of D-Health. In the rest of the paper, we provide a technical feasibility study for these technical components. A technical feasibility study for challenge (v) is provided in Section 5, followed by technical feasibility studies for challenges (vi) and (vii) in Sections 6 and 7, respectively.

2. PROPOSED SYSTEM ARCHITECTURE

We define a remote health monitoring and management system as a *system* that provides the interface between a *patient* and a *doctor*, as shown in Fig. 1. The system acquires, stores, and analyzes patient health data along this transition. Although a much finer grain sub-layering of a typical remote health monitoring system is possible, our proposed system consists of two super-layers: *Front End* and *Back End*. These two super-layers contain similar technical functionality and business opportunities, hence our rationale for this layering. Details of each layer are provided in the rest of this section. Section 2.1 details the Front End, which is the interface between “the patient” and “the system.” Section 2.2 details the Back End, which is the interface between “the system” and “the doctor.”

2.1 Front End

The front end of the system is responsible for acquiring healthcare data from the patient and transmitting it to the back end securely and in a privacy-preserving fashion. There are well-established standards for the acquisition of health data, such as ISO/IEEE 11703-20601:2010 (Fernandez & Pallis, 2014). The connection of this layer to the back end is usually through the internet (Hassanalieragh, et al., 2015), making it necessary to ensure data privacy during acquisition and transmission. The functions of the front end are detailed in this subsection.

IoT-based Acquisition infrastructure: Although the IoT concept is in its infancy, a particular radio communication technology to improve the active bandwidth by deflecting IoT traffic from the internet through a special ultra-high-bandwidth and energy efficient cellular network (900 MHz) has been created by the French company SIGFOX (SIGFOX, 2016). The SIGFOX IoT network will be first deployed in San Francisco. For general IoT networks, three widely available wireless technologies are: i) 3G/4G cellular wide area networks, ii) Wi-Fi local area networks, and iii) Bluetooth Smart personal area networks. A dedicated IoT network has also been proposed as a research topic (Fernandez & Pallis, 2014).

Privacy of the acquired data: In addition to assuring data privacy at a cryptographic and system level (Kocabas & Soyata, 2016; Kocabas & Soyata, 2015), security concerns arising from sensor tampering (Page, et al., 2015b) and sensor data trustworthiness (Kantarci & Mouftah, 2014; Pouryazdan, Kantarci, Soyata, & Song, 2016) must be taken into account in this layer. To create a secure overall system, an *adversary model* must be defined. The most common

adversary model that we will adopt is the *honest but curious* adversary model (Cao, Wang, Li, Ren, & Lou, 2014; Goldreich, 2004), in which a given part of the system is assumed to perform its duties correctly (i.e., “honestly”), but is capable of intentionally or unintentionally observing other parties’ data (i.e., “curious”). Such a system is also vulnerable to side channel attacks from parties that observe the peripherals (i.e., “sides”) of the system and attempt to guess the underlying data. Among the many, a few examples of these attacks include power analysis attacks (Kocher, Jaffe, & Jun, 1999), timing attacks (Kocher P. C., 1996), fault-based attacks (Boneh, DeMillo, & Lipton, 1997), and cache attacks (Bernstein, 2005).

Preprocessing of the acquired data: The amount of the acquired data could be unmanageable in terms of storage, transmission, or processing. Therefore, it is necessary to apply preprocessing algorithms to the acquired data to reduce its size (Soyata T. , Muraleedharan, Funai, Kwon, & Heinzelman, 2012; Soyata T. , et al., 2012). These algorithms are applied to a set of aggregated data, rather than the raw data. The hardware components that aggregate the data from the IoT-based sensors are *concentrators* (Zhao, Wang, & Nakahira, 2011; Hu, Xie, & Shen, 2013). The purpose of a concentrator is to reduce the power consumption of the individual IoT devices by directly receiving the sensor data from them at a short distance and transmitting the aggregated data over much longer distances. While this data concentration is a much higher workload than what the individual IoT devices can handle, concentrators are not necessarily the destinations where the pre-processing takes place. For pre-processing, *cloudlets* are used that are substantially more computationally capable than concentrators and have dedicated WAN links (Soyata, Ba, Heinzelman, Kwon, & Shi, 2013; Powers, et al., 2015). The pre-processing of the data turns raw data into a much more summarized format, such as the computation of the QT and heart rate information from raw ECG signals (Page, et al., 2015c).

2.2 Back End

The back end of the system is responsible for storing and processing the data securely. The functions of the back end are detailed in this subsection.

Secure Storage: The data is acquired in a time series fashion. To store, retrieve, and query time series data, REST APIs are provided within (Zhang, et al., 2013). One of the concerns about handling the data in the cloud is identifying the attack patterns. One example solution, Zachman Framework for enterprise architecture modeling, identifies attacks patterns by checking six characteristics (who, what, where, when, why and how). The patterns of access in the cloud are compared against an independently-running “plane” to determine whether each access is normal or malicious (Blackwell & Zhu, 2014).

Secure Computation: While static storage of data is feasible by using well-known secure storage standards such

as SSAE16, this data cannot be operated on. If computation has to be performed on the data that is stored in an untrusted cloud, emerging cryptographic mechanisms such as Fully Homomorphic Encryption (FHE) are required (Kocabas O. , et al., 2013). These algorithms allow the cloud to perform “blind-folded computation” without observing the underlying medical data, thereby eliminating concerns regarding data privacy (Page, Kocabas, Soyata, Aktas, & Couderc, 2014), however, computations using FHE are orders of magnitude slower than their AES-based “traditional” cryptographic counterpart (Page, Kocabas, Soyata, Aktas, & Couderc, 2014; Kocabas & Soyata, 2015).

Database Sharing: Much like the concentrator in the front end, a portion of the back end is responsible for aggregating databases and sharing them across many applications or other clouds. The key element of this functionality is to aggregate the databases in an identity-removed fashion. Data obfuscation and identity removal are well-established techniques (IBM, 2016) that obfuscate the data in a way that makes the data un-identifiable even if compromised. This functionality of the back end is important since the accuracy of the analytics engine improves as the database sizes grow, thereby improving the statistical inference related to disease detection.

Visualization: The visualization engine can be thought of as being the “visual aggregator.” This engine turns an enormous amount of data into a format that is easily comprehensible and understandable by a human, i.e., the doctor. Despite occupying mega- or gigabytes of storage, the information content in the acquired raw data is very low. The visualization engine is necessary to turn the raw data into a highly summarized format, potentially occupying many orders-of-magnitude less physical space for the same (or higher) information content.

Analytics Engine: Although strictly visualizing the data in a summarized format allows the doctor to access patient information much faster, this visualized information can be further augmented with machine learning (ML) algorithms. The function of the analytics engine is to run a standard set of machine learning and statistical inference algorithms to determine the likelihood of certain diseases for a given set of acquired data. These statistical inferences can be included in the summarized data provided by the visualization engine. The inferences provided by the ML algorithms are much simpler than the visualized data. For example, while a 24-hour visualization (plot) of the patient ECG information could provide the doctor extremely useful and summarized synopsis of the patient’s heart condition, it is still a lot of data to browse through. This plot could be augmented with a single statistically-inferred value (e.g., 87% probability that the patient has the LQT1 heart condition (Page, et al., 2015c). While the initial plot allows the doctor to use his/her experience and knowledge to potentially reach the same decision, augmenting the plot with such a suggestion provides at least a “machine-based second opinion” to the doctor. In the best case, it provides a

“good starting point” or even “the solution that was not obvious to the doctor initially, but was mathematically the best inference.”

3. BUSINESS OPPORTUNITIES

In this section, we will identify the business opportunities in the front end and back end layers.

3.1 Categorizing the Business Opportunities

While a third party business entity can offer the entirety of the services encompassing our system in Fig. 1 as a remote health management and monitoring service, separating the front end from the back end makes sense due to the major characteristic differences that each layer represents. Based on the structure we introduced in Section 2, the front end of the system can be thought of as the “data acquisition” layer, while the back end can be thought of as the “data handling layer.”

Acquisition of the data implies a direct physical connection to the patient; This “physical” connotation significantly limits the location of the third party companies that can provide these services. Alternatively, the back end services could be completely “virtual,” since the offered services are generally “software” in nature. In the following subsections, we provide a detailed list of the potential services in each layer.

3.2 Front End Business Opportunities

In this subsection, we will identify the business opportunities related to the front end. Due to the “physical” nature of the front end, most of the services that can be offered involves making a physical contact with patient at some point.

IoT Hardware and Communications: IEEE standards form a basis for common wireless technologies which are the main component of the front end layer. Relevant sensor networks as LANs are: Wi-Fi 2.4 GHz & 5 GHz (IEEE 802.11ac); and low-power 900MHz (IEEE 802.11ah); ZigBee & ZigBee PRO 2.4 GHz & 900 MHz (IEEE 802.15.4) and 6LoWPAN (for IoT); as PANs are: Bluetooth 2.4 GHz (IEEE 802.15.1) and Bluetooth Low Energy (BLE); UWB (IEEE 802.15.4a); RFID (IEEE 802.15.4f) (Soyata, Copeland, & Heinzelman, 2016) and Low Rate WPAN (IEEE 802.15.6), which are identified by the IEEE standards for the body area network (BAN). As very low-power: DASH7 based on 433 MHz (ISO/IEC 18000-7). GSM, GPRS, UMTS, HSPA and LTE are the current standards in mobile cellular networks (Fernandez & Pallis, 2014). Wired and wireless communications working on same infrastructure will concentrate with the 5th generation communication technology (5G) for people and IoT. The future networked society will run on this omnipresent communication technology which has ultra-high bandwidth (EC Horizon 2020, n.d.). Enabling breakthrough user controlled privacy, wireless connection to

over 7 trillion devices for over 7 billion people, better optimization for storage, processing and big data analytics, 90% energy saving per service and 100 times higher wireless bandwidth compared to 2010 are expected to be allowed by the 5G technology (Fernandez & Pallis, 2014). These technologies are considered to support the communication between the devices used for health monitoring. Thus, *network & communication* providers such as Verizon, AT&T or Cisco are expected to serve with one or more aforementioned wireless communication technologies in this layer.

Sharing (renting) databases: The necessity of physical contact to the patient doesn’t necessarily mean that each contact creates a single business opportunity. Although the initial data must be acquired by making a physical contact to the patient, say, patient A, this data can be used to provide a data sample for patient B through database sharing. As described in Section 1, data analytics algorithms will work more accurately when the information from patient A+B is available, as opposed to only patient A. Therefore, this creates a business opportunity for the company that acquired the data from patient A. With proper user consent, the third party can anonymize the database using the obfuscation software described in Section 2.2. The obfuscated data provides a business opportunity to be “rented” to other third parties, or, corporations such as the insurance companies for use in data analytics.

Self Data Acquisition: One of the important implications of technology is that users do not have to have deep knowledge of the inner workings of the devices to use them. Data acquisition for routine monitoring tasks, such as personal ECG monitoring, can be done without the intervention of a healthcare professional. However, this doesn’t mean that no business opportunity exists for simple data acquisition tasks like this. Smartphone applications that are approved by healthcare organizations can be sold to allow users to acquire their own data. The purpose of the smartphone app is to significantly simplify the user’s job by providing visual instructions, whether static or interactive and calibrate the sensors by directing the user through multiple steps. These smartphone apps can be sold either with the sensors through pharmacies, or separately.

Professionally-Assisted Data Acquisition: When the level of complexity to acquire the data exceeds a level that no longer allows a simple smartphone app to be used, healthcare companies specialized in data acquisition can sell their services to acquire the medical data. This can involve bringing concentrators, cloudlets, and sensors to the user’s home and attaching them in proper order and ensuring proper communication with the cloud. In many cases, using the professional services might be legally necessary due to the legal implications involved in the well-being of the individual.

Invasive Data Acquisition: In the extreme case, the data acquisition might require a surgery, such as the implanting of a defibrillator. Clearly, this operation might

be only feasible in a hospital environment or an approved company with such expertise. The important note to make here is that, such a service is in fact a separate component of the overall system and does not necessarily have to be provided by the provider of the rest of the services.

3.3 Back End Business Opportunities

In this subsection, we will identify the business opportunities related to the back end. Since this layer does not represent a “physical” contact with the patient, it can be provided virtually anywhere.

Infrastructure as a Service: The infrastructure to store and manipulate medical data can be rented through the widely-accepted Infrastructure as a Service (IaaS) concept. Rather than a generalized infrastructure, a more specialized infrastructure provides much better business opportunities (Powers & Soyata, 2015). For example, the databases that store medical information could be optimized to handle medical data, potentially incorporating privacy preserving storage and data obfuscation methods as built-in features. Companies such as IBM, Oracle, Microsoft and Teradata are the potential service and technology providers for this business opportunity.

Disease detection (Analytics) algorithms: Although well-known standard algorithms exist for detecting certain diseases, a one-size-fits-all algorithm is not possible due to the sophisticated biological processes involved in different diseases. Therefore, a new algorithm that achieves a higher detection rate using the same database could provide a significant business opportunity to a healthcare organization that wants to use it for patient monitoring.

Visualization Algorithms: Visualization algorithms can be thought of being a sub-category of Software as a Service (SaaS). As will be exemplified in Section 6, the only difference from SaaS is that the visualized data could be displayed with either static limits, which do not depend on a database, or dynamic limits, which do. In the specific case of ECG visualization that we will show in Section 6, the knowledge of the specific disease that is being displayed is crucial. Therefore, the provider of the visualization services is not just renting the software, but the database and disease expertise too. So, it is highly likely that, depending on the disease that is being visualized, the visualization algorithms and their operation change dramatically.

Prediction and Analytics Services: In addition to providing the algorithms as a service, the results of the algorithms also provide an opportunity to rent as statistics in certain diseases. Parties interested in such information are organizations like CDC, or insurance companies that want to compare the disease occurrence rates in certain geographical regions.

4. BACKGROUND AND RELATED WORK

In this section, we will introduce the technical details of different sub-layers. In the following sections, we will

perform a technical study of some of these layers. A three tier architecture can be considered for most proposed frameworks in terms of health monitoring as follows: 1) Wireless Body Area Network (WBAN) for wearable sensors to gather the data, 2) Communication and networking, and 3) The service layer (Pantelopoulos & Bourbakis, 2010; Paradiso, Loriga, & Taccini, 2005; Milenkovi, Otto, & Jovanov, 2006; Bazzani, Conzon, Scalera, Spirito, & Trainito, 2012; Benharref & Serhani, 2014). Various physiological parameters such as blood pressure and body temperature can be measured by the wearable sensors as proposed as a system model in (Babu, Chandini, Lavanya, Ganapathy, & Vaidehi, 2013). A Bluetooth connection is used for relaying the acquired information to a gateway server by sensors. The gateway server converts the data to an Observation and Measurement file and keeps it on a remote server to be acquired by clinicians via internet. In (Rolim, et al., 2010), a health monitoring system is presented to illustrate medical staff reaching the stored data online through content service implementation to utilize a similar cloud based medical data storage. To supervise patients with high risk of heart failure, WANDA (Lan, et al., 2012), an end to end remote health monitoring and analytics system is presented by aiming at a particular medical implementation.

4.1 Data Acquisition and Sensing

Wearable devices which combine a communications platform to convey the measured data, hardware for minor preprocessing and miniature sensors that measure various physiological parameters acquire physiological data. Wearable sensors that are or will be available to measure some biomarkers encapsulated in Table I. Those biomarkers that can diagnose four common disease categories have the applicability levels which are also indicated in the table.

Wearable sensors have some physical limitations due to wearability requirements which are being lightweight and small and not blocking patients’ maneuverability. Moreover, energy efficiency is essential for those because of the limited place for the batteries in the wearable package.

TABLE I

A list of advanced sensors and their potential application potential to the monitoring of certain diseases. ** indicates excellent application potential, while * indicates some potential for application.

Biomarker	CVD	COPD	Parkinson’s
Gait (posture)	**	**	**
ECG	**	**	*
Respiratory rate	**	**	*
EMG	*	*	*
Blood pressure	**	*	*
Title volume	**	**	*
Body movements	*	*	**

Highly durable batteries are eminently preferable to provide convenience and to ensure that data is not lost during recharging or replacement despite the rechargeable or replaceable features.

A challenge for the quality of the data captured in terms of the achievable signal to noise ratio can also be presented by the energy efficiency requirements. The closer contact with the skin enables measurement of relatively more physiological parameters and with better accuracy, therefore recent flexible sensor designs (Son, et al., 2014; Xu, et al., 2014; Kim, Ghaffari, Lu, & Rogers, 2012) that can be located in contact with the skin in various body parts are especially alluring for medical implementations. Additionally, there have been attempts to make the operational lifetime of wearable sensors longer by combining low power device and circuit level techniques (Olorode & Nourani, 2014; Park, Chou, Bai, Matthews, & Hibbs, 2006) and energy harvesting method (Torfs, Leonov, Van Hoof, & Gyselinckx, 2006). Furthermore, the operational durability can be improved more by utilizing smart sensing methods on system level.

There have been studies about energy efficient sensing mechanisms in the related background of wireless sensor networks (WSNs) that are accustomed to sense physical phenomenon in a distributed fashion. Current WSNs methods can be referred again to fulfill our requirements despite the more concentrated sensor deployment, compared to WSNs, in our health monitoring system. The suggested energy efficient sensing approaches hinge on appointing sensing duties to the nodes formed on their relative distance so as to sense the maximum amount of physical information while improving energy efficiency by abolishing probable unnecessary sensing duties (Madhani, Tauil, & Zhang, 2005; Chou, Rana, & Hu, 2009) and by distribution of duties formed on the energy availability at each sensor (Zhang & Hou, 2005; Huang & Tseng, 2005; Chen & Zhao, 2005; Cardei & Wu, 2006; Yu & Sharma, 2010). By construing and operating an active context which is formed on energy availability and the patients' health condition, our system can employ akin mechanisms. For instance, as demonstrated in Table I, separately sensed biomarkers have distinctive levels of applicability for particular health conditions. The other sensors are turned off to enable lifetime extension when the energy level is critical and the patients sensitive condition forces focusing on a specific biomarker. The application of such schemes to develop energy efficiency adaptively by approving dynamic utilization of sensors formed on the context is enabled by an IoT based sensing architecture. It is hard to find such flexibility and intelligence in the ordinary data acquisition systems where the gathered information is transmitted passively by sensors. More sophisticated algorithms can also be implemented without patients' manual intervention to wield the sensors or the software on the data concentrator by removing the decision making process to sense task assignment to the cloud.

As the communication can account an important part of the energy usage in sensing devices, appropriate low power communication protocols usage is constrained by energy limitation of these devices. In order to support communication between low power devices that perform in personal operating space (POS) of roughly 10m, ZigBee over IEEE 802.15.4 is usually used in low rate WPANs (LR-WPANs) (Lee, Su, & Shen, 2007). Energy efficient dependable mesh networking is provided by ZigBee. Another wireless communication protocol, Bluetooth low energy (BLE), is appropriate for low power short range communication which is advisable for the particular necessities of implementations such as health monitoring, sports, and home entertainment. The design purpose of the original Bluetooth protocol (IEEE 802.15.1) is to provide relatively short range communications for implementations of a streaming nature, such as audio. By applying extended sleep intervals to enable the general energy efficiency, the framework is altered by BLE. A superior energy efficiency in terms of number of bytes sent per Joule of energy is accomplished by BLE (Siekkinen, Hienkari, Nurminen, & Nieminen, 2012). An intermediate node (data concentrator) is required to make sensors data and control available over internet while the preceding communication protocols in use. IPv6 through Low Power Wireless Personal Area Networks (6LoWPAN) has been put forward to perfectly connect energy constrained WPAN devices to the internet to comprehend the IoT concept additionally (Bui & Zorzi, 2011). In order to fit IPv6 datagrams into IEEE 802.15.4 limited frame size to enable IP access to low power, low intricacy sensing devices, fragmentation techniques are construed by 6LoWPAN.

4.2 Internet of Things (IoT)

Integration of the IoT paradigm with electronic remote health monitoring systems can boost flexibility, intelligence and interoperability more while being noticeable of these systems has assured to transform the conventional health care methods (Bazzani, Conzon, Scalera, Spirito, & Trainito, 2012) (Ray, 2014). With the help of IoT architecture, identifiable and uniquely addressed devices are available through the internet at anytime and anywhere. Devices equipped with IoT functionality in health monitoring systems, fairly reducing the work load on set up and administration tasks, can exchange information with each other and health institutes in addition to their capability of conventional sensing tasks. Providing services such as automatic alarm to the closest healthcare institute during the time of a critical accident for a supervised patient can be an example for such systems (Bui & Zorzi, 2011).

4.3 Cloud Storage and Processing

Most of the research on sensors related to healthcare monitoring deal with managing of the data on the devices, storing the medical data directly on computer nodes, or

utilizing intermediate nodes for storage and/or computation. Data storage and management through the Cloud Technology has been pointed rarely pointed out in the related work. A sensor-oriented cloud infrastructure is presented by the authors in (Yuriyama & Kushida, 2010). The actual devices are not included by the early evaluation results since the initial results are formed on simulated sensors in the preceding principle. To store sensor-based data in a dedicated manner, nonetheless, several Cloud-based services are currently available (e.g., Pachube, Nimbits, ThingSpeak, iDigi). Available services call for solutions for data security and provisioning interface for linkage to mobile or external implementation on latter processes (Doukas & Maglogiannis, 2012).

Cloud Computing enables favorable, on-demand network access to adjustable computing resources configured as a shared group which an interaction by service provider or slightest managing attempt is enough to provision or release in a quick response time. Devices, like smart phones, considered as heterogeneous thin or thick platforms can reach resources through the network accessing over standard mechanisms. Virtual machines, network bandwidth, memory processing, and storage are mentioned to be examples of resources. The dexterity which develops with users being able to quickly and cheaply re-provision technological infrastructure resources is a huge asset by given the essences of Cloud Computing and the resilience of the services which can be improved. Since there is no need for a specific location or device, any user can connect the system using a web browser in any location and on any device. Allowing for unification infrastructure in locations inexpensively is possible due to multi-tenancy facilitating resource and costs sharing by virtue of enormous user pool. To manage user data, many Cloud Computing applications are accessible for both free (e.g., iCloud, Okeanos, Pithos, Dropbox) and commercial usage (e.g., GoGrid, Amazon AWS, Rackspace). Building custom applications and consolidating Cloud Computing functionality are not supported by many of them. Furthermore, optimization to service healthcare-based implementation is not validated yet (Doukas & Maglogiannis, 2012).

4.4 Analytics and Visualization

Medical data analysis and visualization are also critical elements of remote health monitoring implementations (aside from the technology for data acquisition, storage and access). It is essential to analyze the medical records containing various physiological characteristics over a long period of time to diagnose accurately and to monitor a patient's medical condition. The data analysis task becomes frustrating and error prone for clinicians who work with multidimensional data, especially when long term (i.e. high quantity) of data is used. Data mining and visualization techniques have just reached the considerable level in

remote health monitoring implementations (Ukis, Tirunellai Rajamani, Balachandran, & Friese, 2013; Rao, 2013), despite the fact that they have been addressed as a solution to the preceding challenge (Wei, et al., 2005; Mao, et al., 2011).

We have proposed that decision support be performed by a dedicated company, which may or may not be responsible for collecting the data. This "Clinical Decision Support as a Service" has also been described in (Weaver, Ball, Kim, & Kiel, 2015), where they suggest standardizing the relevant portions of healthcare records to make computerized analysis easier. (Halamka, 2010) agrees with this, and further suggests standardizing the decision support rules, which we believe will be difficult when using machine learning and in the presence of competition. An "open" version of the cloud-based health monitoring concept is discussed in (Li, Guo, & Guo, 2014), in which scientists and healthcare professionals can share their data and models. We expect that such a system may still thrive alongside paid solutions, while the proprietary versions may be based on specialized databases and more refined algorithms.

5. FRONT-END FEASIBILITY STUDY

At this point, we have described all of the pieces of our ideal remote monitoring system. We will now present case studies which detail specific components.

The front end segment of remote health monitoring is anticipated to be connected to the Internet of Things architecture as this segment is responsible for data acquisition through wearable sensors, and the sensors are to be interfaced via front-end circuitry in nearby devices that would offer built-in IoT sensing capability. These include smart phones, smart tablets and other personalized devices with communication interfaces. A typical cloud-inspired service model to acquire data through the front-end of the




Figure 2. Front-end design by IoT sensors interfacing wearables.

health monitoring could be Sensing as a Service (S²aaS) (Sheng, Tang, Xiao, & Xue, 2013). In a cloud-centric IoT architecture, uniquely identifiable sensors push data to the cloud platform for being aggregated, analyzed and presented to the end user (Gubbi, Buyya, Marusic, & Palaniswami, 2013). Data acquisition through cloud-centric IoT has to maximize the usefulness of the collected data for the platform whereas the sensing costs of the IoT sensors may need to be compensated. A minimalist illustration of the concept can be seen in Fig. 2, and the following mathematical model can be used to analyze the feasibility of such front-end structure.

Utility of the cloud platform can be calculated as the difference between the total usefulness of the data and the compensation made to the IoT sensors for their sensing costs in a certain time window. It is worthwhile noting that we use data to denote a task of sensing a particular phenomenon. In (1), U_p denotes the utility of the cloud platform while $V_\tau(S_\tau)$ stands for the usefulness/value of the data received for sensing tasks handled by the IoT sensors in the set, S_τ during the time window, τ . In the same equation, ρ_s^τ denotes the sensing cost/compensation of sensor s of the overall sensors set, S during the time window, τ .

$$U_p = \sum_{\tau} \left(\sum_{t \in T_{S_\tau}} v^\tau(S_\tau) - \sum_s \rho_s^\tau \right) \quad (1)$$

Besides the utility of the platform, utility of the nearby IoT sensors is another metric that is to be used in the feasibility study of the front-end segment in a remote health monitoring system. If the IoT sensor is compensated based on the usefulness of its sensor reading, the compensation should be no less than the sensing cost. Equation (2) formulates the utility of an IoT sensor (U_s) as the difference between the total compensation received for participating in the sensing tasks and total sensing cost.

$$U_s = \left(\sum_{\tau} \left(\sum_i \rho_i^\tau - \sum_i c_i^\tau \right) / |S_\tau| \right) / \tau_{end} \quad (2)$$

As the wearable sensors are interfaced with nearby mobile devices and their corresponding built-in IoT sensors, the cloud platform can be misinformed due to either of the following scenarios: 1) Built-in sensors of mobile devices may be malfunctioning. 2) Users may be behaving maliciously to send altered sensor data. Regardless of the intention of the IoT sensor, misinformation/disinformation of the cloud platform may lead to severe consequences in patient's health. In other words, platform utility is significantly reduced if wrong sensor data is shared with the cloud platform. Here, trustworthiness of the IoT sensors introduces an important consequence impacting the platform utility. Reputation-based models can be utilized to reduce manipulation probability in the aggregated data at the cloud

platform. In (Kantarci & Mouftah, 2014a; Kantarci & Mouftah, 2014b; Kantarci & Mouftah, 2014c), trustworthy data acquisition schemes have been proposed for public safety purposes in a cloud-centric IoT architecture. This concept can easily be adopted by the front-end segment of the presented remote health monitoring architecture in this paper. In case a particular task is sensed by multiple IoT sensors, the percentage of positive readings upon detection of outliers can be used via an outlier detection algorithm (Zhang, Meratnia, & Havinga, 2010). The statistical reputation of an IoT sensor (i.e., sensor i) at the end of the time window t ($R_i(t)$), can be formulated as shown in (3) where $p(t)$ and $n(t)$ denote positive and negative readings, respectively. Thus, instantaneous reputation and previous reputation undergo a weighted sum function, and an IoT sensor with low reputation will be less likely to be selected and vice versa. Moreover, the usefulness of the data provided by an IoT sensor will be scaled by the reputation of the sensor.

$$\mathfrak{R}_i(t) = \alpha \cdot \frac{p(t) + \epsilon}{p(t) + n(t) + \epsilon} + (1 - \alpha) \cdot \mathfrak{R}_i(t - 1) \quad (3)$$

Adopting the IoT-based data acquisition in the front-end can increase the precision of sensed data as the higher the number of sensors the better the performance of a distributed estimation system. On the other hand, due to the issues reported above, trustworthiness of the data acquired through IoT sensors can be guaranteed by reputation-based sensing. Fig. 3 illustrates the disinformation probability in a distributed sensing environment under reputation-unaware sensing and reputation-aware sensing in the presence of malicious behavior and malfunctioning sensors where sensing costs of IoT sensors vary between 1 and 5 and the usefulness of sensor data varies between 1 and 10. In the experimental setup, 1000 IoT sensors are deployed in a




Figure 3. Manipulation probability in the presence of malfunctioning and malicious IoT sensors.




Figure 4. (a) Utility of the cloud platform in a distributed IoT sensing scenario, (b) Average utility of an IoT sensor node in a distributed IoT sensing scenario.

1000x1000 terrain with 5% malfunctioning or malicious activity. It is worth noting that disinformation denotes the case where an IoT sensor is recruited while it is reporting wrong sensor data. Reporting of wrong sensor reading can be either continuous or intermittent. Intermittent disinformation/misinformation denotes the situation where true sensor data is sent for a while and then wrong sensor data is shared either due to malfunctioning or to lead to disinformation. As seen in the figure, reputation awareness degrades disinformation probability at the order of 75% under various sensing task arrival rates. Malicious or malfunctioning sensors can be identified faster if they keep sending wrong sensor data continuously. Thus, the reputation of a sensor that continuously sends wrong sensor reading will be degraded continuously and converge to a low value shortly, and the corresponding IoT sensor device will not be recruited again due to its reasonably low reputation. Therefore, disinformation probability under intermittent disinformation is slightly higher, however, the improvement over reputation-unaware sensing is still above 70% even in the presence of malicious sensors that attack based on a strategy.

In addition to the experimental results above, Fig. 4 presents the utility of the cloud platform and the average utility of an IoT sensor calculated by (1)-(2) in the presence of malfunctioning and malicious IoT sensors which may report wrong sensor data either continuously or intermittently. The simulation setup is adopted from the study in (Kantarci & Mouftah, 2014a). As seen in Fig. 4a, platform utility can be improved by 12% under a lightly arriving sensing task load and by 85% under a heavily arriving sensing task mode. As seen in Fig. 4b, compensation of IoT sensors is always non-zero. Note that the compensation mechanism in these examples adopts the auction-based payment approach in (Yang, Xue, Fang, & Tang, 2015).

6. VISUALIZATION CASE STUDY

In order to concretely illustrate backend components, Sections 6 and 7 will focus on a single case study: detection and monitoring of the Long QT Syndrome (LQTS).

6.1 Background: Long-QT Syndrome

LQTS is a cardiac illness which may be congenital or drug-induced. It is characterized by prolongation of the QT interval on an ECG, shown in Fig. 5. This interval is a measure of ventricular repolarization time, and its prolongation can warn of impending arrhythmias such as torsades de pointes (TdP), leading to syncope or death. The congenital form of the disease is particularly dangerous, as this risk never fully goes away.

The impact of LQTS varies widely based on gender, age, and specific genetic mutation. It also manifests during different activities based on genotype. Type-1 LQT (LQT1) patients tend to have issues during exercise, while Type-2 (LQT2) patients are more at risk during sleep. Patients with an LQT genotype, or people who are on known QT-prolonging drugs, may benefit from — or outright require — long term monitoring via ECG sensors, providing an early warning to the patient, doctor, and/or EMS based on QT interval measurements. More specifically, the physician is interested in the length of the QT interval in relation to heart rate; i.e., whether QT is happening quickly enough, before the next cardiac cycle begins. It is common to look at a “corrected” QT value, shown in Equation (4), known as the Bazett QT correction equation (Bazett, 1920). While it is not necessarily the best correction for all purposes, it is perhaps the simplest and one of the most commonplace.

$$QTc = \frac{QT}{\sqrt{RR/sec}} \quad (4)$$




Figure 5. Standard ECG waveform. We will mainly look at the QT interval (annotated), but RR — the time of one heart beat — is also of interest. Other metrics may provide even more detail, such as Tpeak–Tend (TpTe). Image source: SinusRhythmLabels.png by Anthony Atkielski.




Figure 6. Example plot of QTc over 24 hours in the “ECG Clock” format. This patient has a relatively normal QTc interval during the day, but it becomes potentially dangerous at night.

6.2 Decision Support

One of the most useful types of decision support is not for a computer to generate specific recommendations, but to simply present the data in a manner that allows the doctor to fully understand the situation. Based on this presentation, the doctor can make his or her own decision. The challenge here is to condense many sensor measurements spanning a long period of time into a very concise summary.

An important consideration in building visual aids for decision support is knowing which features are relevant to the condition being investigated. In the Long QT Syndrome (LQTS), for example, many ECG measurements such as QT, RR, or TpTe may all carry some information about the disease, not just QT. Additionally, there are several ECG leads (sensor locations) to choose from, and certain leads may be better for QT measurement. We also know that LQTS manifests differently throughout the day based on patient genotype, so perhaps there are a few key times of day that should be checked (as opposed to looking at an overall average of the available data).

We are building a sizable array of factors that are relevant to this disease, and circling back around to the original problem: displaying it all to the doctor in a form that can be digested very quickly. Remember that in addition to each of these factors — ECG marker and lead, time of day, etc. — the doctor may also have 20–30 patients. Further, the advent of long-term remote monitoring means that each patient will be generating more data than ever before. So we would like to create a picture that adequately summarizes a patient’s day with only a few seconds of viewing (Page, et al., 2015c).

The first set of techniques we will apply to LQTS monitoring involve the removal of redundant information from the ECG recording. For instance, while many ECG measurements may contain some information related to the patient’s illness, we may focus simply on QTc (which combines two measurements, QT and RR). Further, since many ECG leads are available, we will combine data from all of them using e.g. a median or average (We could also choose to look only at a single lead, perhaps the least noisy.)

Now that we are focused on a single (computed) feature on a single (virtual) lead, our visualization problem is much more focused. We must plot or tabulate the values of QTc for ~120,000 heart beats per day. Again, we know that certain times of day are more critical based on genotype. However, as they are based on sleep and exercise patterns, they will still vary significantly between patients. So, we would like to show the entire day if possible.

The most obvious way to present the remaining data would be to simply plot it. However, the scale of the plot must be determined to ensure that short duration events are still visible. In the case of LQTS, we are mainly interested in events that last for several minutes. It is therefore practical to plot a full 24 hours in a fairly “typical” plot size (e.g. “half page”), which allows us to see with at least one-minute resolution.

Finally, we note that for data spanning 24 hours or more, polar axes can be beneficial. By using the angle of a polar plot to represent time of day, and the radius to represent the value of some feature, multi-day data can simply continue to circle around the plot. Even with single-day data, this representation makes it unnecessary to adjust

the axes range to view different recordings. (For example, should the x axis start where the recording does, or at some other time like midnight?) We have found it best to standardize on a 24-hour polar axis. An example of the visualization we've just described is given in Fig. 6. Because our data was still fairly noisy even after all the preceding steps, we used a median filter to smooth it. Further uses of this visualization technique have already been well-described in (Page, Soyata, Couderc, & Aktas, 2016) and an open source ECG visualization program is available in (Page, Soyata, Couderc, & Aktas, 2015a).

7. ANALYTICS CASE STUDY

The objective of the visualization techniques developed in Section 6 was to present enough data for the doctor to make a decision. However, especially in the case of rare diseases with which the physician may not be experienced, it would be good for the computer to also provide some extra “hints”. In this section, we begin to investigate ways to augment the visualizations using machine learning (ML) algorithms. The goal is to utilize (a subset of) the same data used to generate plots to compute the likelihood that a patient has a particular medical condition. In this study, we will continue to focus on LQTS.

7.1 Background/Methods

We consider machine learning algorithms from three general categories:

- 1) “Conventional” supervised learning methods, such as SVM, decision tree, and nearest neighbors.
- 2) Clustering techniques such as GMM, K means, and DBSCAN.
- 3) Artificial neural networks (ANN).

We will mainly discuss the first category, but will present some formulation and results from the third. We will also consider “ensemble” techniques such as AdaBoost and Random Forest, which attempt to use the results from several classifiers improve accuracy. Clustering methods will not be discussed, as we have not yet found satisfactory parameters to achieve good results with these.

Which ML algorithm is best to detect and classify LQTS? This really depends on properties of our data and our long term goals for how it will be used. For instance, some algorithms may be lighter in terms of storage and/or computation if we intend to continuously update the classifier (i.e. “online” machine learning). Additionally, we will want to keep the dimensionality of the data as low as possible in order to improve the accuracy of many methods. For now, we will make some assumptions — e.g. that hourly data will be sufficient, as opposed to beat-to-beat data, and that the database is small — and test the performance of a variety of conventional ML algorithms on our data. Incidentally, our database is indeed somewhat

small; we have access to 639 24-hour Holter recordings of healthy, LQT1, and LQT2 patients. LQT2 is the smallest cohort, with 145 recordings. LQT1 has 294 recordings, and healthy has 200. The scikit-learn (Pedregosa, 2011) Python library will be used to perform the tests. 70% of the samples will be used for training, and the remaining 30% will be used for testing. Again, the data will consist of hourly QTc values for each patient, plus their gender (25 “dimensions”), and a corresponding classification (0, 1, or 2, for “healthy”, “LQT1”, or “LQT2”, respectively).

To start with, we will test one of the simplest machine learning algorithms: **nearest neighbors**. This method simply selects the “closest” training sample to the presented sample (i.e. shortest Minkowski distance). An extension of this takes a weighted average of the N closest samples. One disadvantage of this technique is that you must store and search through all previous data in order to find the nearest match(es).

Support vector machines (SVM) are also very common, and simple to train and interpret. Depending on the nature of the data, they can be highly accurate. They operate by defining hyperplanes which separate the data into different groups. These planes are created from a subset of the training points, known as support vectors, in a way that maximizes the distance from the plane to the nearest data point of any class. Additionally, the feature space may be transformed using different kernels to allow nonlinear classification boundaries. Regardless of the kernel, SVM offers several advantages including memory efficiency and effective classification in high dimensional spaces. We will attempt to train SVMs using both linear and radial basis kernel functions. While this (and some other) algorithm(s) are designed for data from only two categories, the scikit-learn implementation will internally split our three-category data into two-category stages to bypass this limitation. One weakness of SVM is in the ability to judge how certain we are about a prediction; you can compute the distance from a point to the nearest separator plane, but this doesn't necessarily translate well to a “confidence percentage”.

A very different method that we expect to perform well is the **random forest algorithm**. This method uses training samples to construct multiple **decision trees** — a forest — using random subsets of the given features to build each tree. It then classifies a new testing point by averaging the results of the individual trees (a single decision tree operates by splitting the data multiple times until “leaves” are created of a single class. Then, to classify a new sample, we simply traverse the tree based on the splitting criteria until we arrive at a leaf.) By taking the mean, a random forest gets rid of the over-fitting problem that is often encountered with a single decision tree. Random forests offer other valuable features such as processing large amounts of data efficiently.

The random forest is a type of “ensemble” algorithm, basing its output on the output of several other classifiers. We will test two other ensemble techniques as well:




Figure 7. Classifying LQTS using QT Clocks and ANNs. Top: typical samples for ANN handwriting classification. Bottom: QTc clocks converted to a similar style. The first three clocks are healthy, followed by three LQTS1 and two LQTS2 clocks. While this format may not be ideal, it allows us to test preconfigured ANNs on our LQTS diagnosis problem.

AdaBoost, and **voting**. The voting classifier simply takes the output of several other classifiers and does a majority vote if they disagree. In a more advanced version of this, the results of the individual classifiers will be weighted based on their confidence in it (and/or our confidence in that classifier). AdaBoost is somewhat different; it is a multistage classifier where each stage is trained on the failures of the previous stage (in our case, each stage is a Decision Tree, but this can be changed).

Finally, we will use the NVIDIA Deep Learning GPU Training System (NVIDIA-DIGITS, 2016) and Caffe (CAFFE, 2016) for ANN-based classification of LQTS. We have seen in Section 6 that proper visual arrangement/presentation of ECG sensor data can greatly aid the doctor's decision in diagnosis and prescription. As there are many ANNs designed for visual recognition tasks, we decided to adapt our visual output (i.e. the ECG Clock) to a form that could be directly used as input for a pre-tuned ANN. One common vision task for ANNs is to classify handwritten digits from the MNIST handwriting database (LeCun, Cortes, & Burges, 1998b).

These are binary images, and are 28x28 px each. We simply shrink our ECG clocks (the plotted lines only) down to this size, and attempt to train an ANN to classify "healthy", "LQTS1", and "LQTS2", from plotted QTc values.




Figure 8. Comparison of conventional ML classifiers. The training+testing cycle was repeated 20 times with Holter recordings randomly assigned as training or testing each time. Here we see the average performance of each classifier when identifying "healthy vs. sick" (blue) or "healthy vs. LQTS1 vs. LQTS2" (red). Error bars show the range from worst- to best case performance of each classifier over the 20 trials.

This format essentially restricts us to 784 data points (28x28), and most of the image is blank (i.e. it is sparse); we may only plot ~70 points. So it will be interesting to see if we are providing enough data to the ANN. This technique is shown in Fig. 7. Based on the examples in this figure, we expect that "healthy vs. sick" will be fairly simple to determine, but "LQTS1 vs. LQTS2" may be difficult.

7.2 Results

Classification of "healthy" vs. "long QT" was relatively accurate — about 90% — as we expected. Further differentiating between LQTS1 and LQTS2 was more difficult, lowering the score of each classifier as shown in Fig. 8. Still, an accuracy of about 70–75% was consistently achievable with the SVM and Random Forest methods. We should note right now that several of the recordings in our database were noisy or incomplete, which likely degraded our results. Missing data was replaced with average values, but very short recordings should probably have simply been thrown out. However, we wanted to present a "worst case" starting point for further research, so all data was retained. Another important consideration is that while our data is segregated by LQTS genotype, not all LQTS subjects show the corresponding phenotype. In other words, a handful of the LQTS patients truly do look healthy, so even a cardiologist would be likely to "misclassify" them.

We just saw that Random Forest and Support Vector Machine (SVM) generally proved superior to other algorithms. Now, we would like to see what information they are using to arrive at their decisions. For example,

based on our findings in (Page, Soyata, Couderc, & Aktas, 2016), we expect that data from ~3AM will be a very good differentiator between the classes. We also expect that afternoon QTc measurements will not help distinguish between LQT1 and LQT2. Fortunately, we can examine the internals of these trained classifiers quite easily. In Fig. 9, we extract the “importance” of each feature (hour) from the Random Forest and Linear SVM classifiers. The results are basically what we expected — late-night data is most useful to both classifiers — but SVM also used information from earlier in the evening. Because of the random nature of the training/testing data split, and the random selection of features in Random Forest, these results will not be exactly the same on every trial. However, we observed the same general trend over several trials.

Finally, we attempted a very basic Artificial Neural Network analysis of the QTc data. In this case, we did not provide hourly data points, but (28x28 px) QTc clocks as shown in Fig. 7 (bottom). These clocks were used to train a LeNet network (LeCun, Bottou, Bengio, & Haffner, 1998a), which is known to perform well on the MNIST handwriting data set. Missing data was not “filled” in any way; we simply passed incomplete plots to the ANN. This implementation achieved ~70% accuracy with absolutely no tuning (and ~90% accuracy when only classifying “healthy vs. sick”). i.e., it was comparable to the classifiers we’ve already discussed. However, providing QTc clocks based on a different correction equation (Fridericia, 1920):

$$QTc = \frac{QT}{\sqrt[3]{RR/sec}} \tag{5}$$

(as opposed to Eqn. 4) yielded a significant improvement: three-way classification accuracy increased to ~80%. Using this alternate QTc equation did not improve the accuracy of any of the other classifiers, only the ANN.

7.3 Future Work

TensorFlow (TENSOR-FLOW, 2016) and Amazon Machine Learning (AMAZON-ML, 2016) are two relatively recent ML developments that we have yet to test. Both of these solutions appear to be relatively simple to use and to collaborate with. The Amazon product is promising as a very high level solution, that will simplify the tuning process. TensorFlow, developed at Google, is likely to be useful for researching and training more complex neural networks.

Another avenue of research will arise as data collection and collaboration increases: the analysis of trends and disease outbreaks. This analysis may be more statistical in nature than what we’ve presented; i.e., machine learning may not play a major role. This will also tie in with the visualizations of Section 6. For example, the statistically “normal” ranges must be continuously updated with new recordings.




Figure 9. Weight of each measurement in classification. A stronger magnitude means QTc measurement was more “helpful” at that time. Random Forest weights are extracted from the classifier’s `feature_importances_array`. SVM weights are taken as the (normalized) maximum amplitudes of the weights in `coef_` across all three possible classes (healthy, LQT1, LQT2). We see that both classifiers focus on late-night QTc values, and that SVM also uses information from evening (~7PM) while Random Forest does not.

In Fig. 9, the low values during the daytime (about 9AM-5PM) indicate that perhaps we don’t even need to monitor that data. In further research, we will attempt to select only the necessary features to provide similar results. If there are indeed several hours which don’t require observation, it may save battery life, storage, and processing time. We must also determine which other features/measurements would improve performance. Gender, for example, was used as a feature in our results above, but its importance turned out to be quite low. We must attempt to add other ECG markers (such as RR and TpTe) to determine the best set of features for classifying LQTS. Other researchers will have to do the same for other diseases. Many optimizations remain in terms of classifier parameters as well, but tweaking them did not affect our performance very much at this stage. We therefore believe that it is more important to find the correct features before finding the optimal classifier configurations.

At this point, our ANN results are really only a very preliminary proof of concept. We must optimize this on two fronts: 1) neural network parameters (layers, etc.) and 2) input data. The input data side will be beneficial to the “conventional” algorithms as well. This is the research we just mentioned, where we will attempt to identify other ECG parameters to include in the input, and how to best reduce the dimensionality of the feature space. Further, we will attempt to hand-select only clean, complete recordings of phenotype-positive individuals as input; from initial testing,

this may reduce our error by an order of magnitude! We may also try a different branch of research, where we look at, for example, 1 hour of ECG data, and attempt to predict if there will be a cardiac event in the following hour. This will allow us to provide real-time warnings, rather than being limited to disease classification.

8. SUMMARY AND CONCLUDING REMARKS

Emerging technologies such as IoT and cloud-based machine learning have opened the door for vast improvements to personalized health care. However, we must understand the strengths and limitations of each technology to assemble a system that is reliable, practical, and provides the best possible support to both doctors and patients. We have addressed many of the privacy and security concerns in this system, and presented our approaches to developing some of the key components. We have also identified several business opportunities that naturally arise from such a system, for instance in the realms of data acquisition, sharing, and analytics. In our front end feasibility study, we discussed IoT-based data acquisition in the presence of malfunctioning/malicious nodes. In our backend feasibility studies, we presented decision support methods for long-term patient monitoring. The first method involved visualization of key features from sensor data, and the second method applied machine learning to these measurements to identify disease states. Our future work will focus on improving the ML-based analysis of long-term medical data.

9. ACKNOWLEDGMENT

This work is supported in part by the National Science Foundation grants CNS-1239423 and CNS-1464273. Tolga Soyata was an Assistant Professor - Research at the University of Rochester during the preparation of this manuscript. He was the PhD adviser of Alex Page, as well as the undergraduate adviser of Shrouq Hijazi. Before the publication of this manuscript he joined SUNY Albany, ECE as an Associate Professor.

10. REFERENCES

- AMAZON-ML. (2016). Amazon Machine Learning. Retrieved from <https://aws.amazon.com/machine-learning/>
- Apple Inc. (2016). Apple watch. Retrieved from <https://www.apple.com/watch/>
- Babu, S., Chandini, M., Lavanya, P., Ganapathy, K., & Vaidehi, V. (2013). Cloud-enabled remote health monitoring system. *ICRTIT*, (pp. 702-707).
- Bazett, H. C. (1920). An Analysis of Time Relations of the Electrocardiogram. *Heart*, 7, 353-370.
- Bazzani, M., Conzon, D., Scalera, A., Spirito, M., & Trainito, C. (2012). Enabling the IoT paradigm in e-health solutions through the VIRTUS middleware. *TrustCom*, (pp. 1954-1959).
- Benharref, A., & Serhani, M. (2014). Novel cloud and SOA-based framework for E-Health monitoring using wireless biosensors. *IEEE Journal of Biomed. and Health Inf.*, 18(1), 46-55.
- Bernstein, D. J. (2005). Cache-timing attacks on AES.
- Blackwell, C., & Zhu, H. (2014). *Cyberpatterns*. Springer International Publishing.
- Boneh, D., DeMillo, R. A., & Lipton, R. J. (1997). On the importance of checking cryptographic protocols for faults. *EUROCRYPT*, (pp. 37-51).
- Bui, N., & Zorzi, M. (2011). Health Care Applications: A Solution Based on the Internet of Things. *ISABEL*, (pp. 131:1-131:5).
- CAFFE. (2016). Berkeley Vision and Learning Center. Retrieved from Caffe deep learning framework: <http://caffe.berkeleyvision.org/>
- Cao, N., Wang, C., Li, M., Ren, K., & Lou, W. (2014). Privacy-preserving multi-keyword ranked search over encrypted cloud data. *IEEE Transactions on Parallel and Distributed Systems*, 25(1), 222-233.
- Cardei, M., & Wu, J. (2006). Energy-efficient coverage problems in wireless ad-hoc sensor networks. *Computer Communications*, 29(4), 413-420.
- Chen, Y., & Zhao, Q. (2005). On the lifetime of wireless sensor networks. *IEEE Commun. Letters*, 9(11), 976-978.
- Chou, C. T., Rana, R., & Hu, W. (2009). Energy efficient information collection in wireless sensor networks using adaptive compressive sensing. *LCN*, (pp. 443-450).
- Dorsey, E. e. (2007). Projected number of people with Parkinson disease in the most populous nations, 2005 through 2030. *Neurology*, 68(5), 384-386.
- Doukas, C., & Maglogiannis, I. (2012). Bringing IoT and Cloud Computing towards Pervasive Healthcare. *IMIS*, (pp. 922-926).
- EC Horizon 2020. (n.d.). Digital Agenda for Europe. Retrieved from <http://ec.europa.eu/digital-agenda/en/towards-5g>
- Fernandez, F., & Pallis, G. C. (2014). Opportunities and challenges of the Internet of Things for healthcare: Systems engineering perspective. 4th International Conference on Wireless Mobile Communication and Healthcare (Mobihealth), (pp. 263-266).
- FitBit Inc. (2016). flex: Wireless activity + sleep wristband. Retrieved from <https://www.fitbit.com/flex>
- Fridericia, L. S. (1920). Die Systolendauer im Elektrokardiogramm bei normalen Menschen und bei Herzkranken. *Acta Medica Scandinavica*, 53, 469-486.
- Goldreich, O. (2004). *Foundations of cryptography: volume 2, basic applications*. Cambridge University Press.
- Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. *Future Generation Computer Systems*, 29(7), 1645-1660.
- Halamka, J. (2010). Decision support service providers. Retrieved from <http://geekdoctor.blogspot.com/2010/06/decision-support-service-providers.html>
- Hassanalieragh, M., Page, A., Soyata, T., Sharma, G., Aktas, M. K., Mateos, G., . . . Andreescu, S. (2015). Health Monitoring and Management Using Internet-of-Things (IoT) Sensing with Cloud-based Processing: Opportunities and Challenges. 2015 IEEE International Conference on Services Computing (SCC), (pp. 285-292). New York, NY.
- Hu, F., Xie, D., & Shen, S. (2013). On the application of the internet of things in the field of medical and health care. *GreenCom/iThings/CPSCoM*, (pp. 2053-2058).
- Huang, C., & Tseng, Y. (2005). The Coverage Problem in a Wireless Sensor Network. *Mobile Networks and Applications*, 10(4), 519-528.

- IBM. (2016). InfoSphere Optim Data Privacy. Retrieved from <http://www-03.ibm.com/software/products/en/infosphere-optim-data-privacy>
- Istepanian, R., Sungoor, A., Faisal, A., & Philip, N. (2011). Internet of m-health Things m-IoT., (pp. 1-3).
- Jawbone, I. (2016). Jawbone fitness trackers. Retrieved from <https://jawbone.com/up/trackers>
- Kantarci, B., & Mouftah, H. (2014a). Trustworthy sensing for public safety in cloud-centric internet of things. *IEEE Internet of Things Journal*, 1(4), 360-368.
- Kantarci, B., & Mouftah, H. (2014b). Mobility-aware trustworthy crowdsourcing in cloud-centric internet of things. *IEEE Symposium on Computers and Communication (ISCC)*, (pp. 1-6).
- Kantarci, B., & Mouftah, H. (2014c). Reputation-based sensing-as-a-service for crowd management over the cloud. *IEEE International Conference on Communications (ICC)*, (pp. 3614-3619).
- Kim, D. H., Ghaffari, R., Lu, N., & Rogers, J. A. (2012). Flexible and stretchable electronics for biointegrated devices. *Annual Review of Biomedical Engineering*, 113-128.
- Kocabas, O., & Soyata, T. (2016). Emerging Security Mechanisms for Medical Cyber Physical Systems. *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)*.
- Kocabas, O., Soyata, T., Couderc, J.-P., Aktas, M., Xia, J., & Huang, M. (2013). Assessment of Cloud-based Health Monitoring using Homomorphic Encryption. *Proceedings of the 31st IEEE International Conference on Computer Design (ICCD)*, (pp. 443-446). Ashville, VA.
- Kocabas, T., & Soyata, T. (2015). Utilizing Homomorphic Encryption to Implement Secure and Private Medical Cloud Computing. *IEEE 8th International Conference on Cloud Computing (CLOUD)*, (pp. 540-547). New York, NY.
- Kocher, P. C. (1996). Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and Other Systems. *CRYPTO*, (pp. 104-113).
- Kocher, P., Jaffe, J., & Jun, B. (1999). Differential power analysis. *CRYPTO*, (pp. 388-397).
- Lan, M., Samy, L., Alshurafa, N., Suh, M., Ghasemzadeh, H., Macabasco-O'Connell, A., & Sarrafzadeh, M. (2012). WANDA: An end-to-end remote health monitoring and analytics system for heart failure patients. *Proceedings of the Conference on Wireless Health*, (pp. 9:1-9:8). New York, NY.
- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998a). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278-2324.
- LeCun, Y., Cortes, C., & Burges, C. J. (1998b). The MINST database of handwritten digits.
- Lee, J. S., Su, Y. W., & Shen, C. C. (2007). A comparative study of wireless protocols: Bluetooth, UWB, ZigBee, and WiFi. *IECON*, (pp. 46-51).
- Li, Y., Guo, L., & Guo, Y. (2014). Enabling health monitoring as a service in the cloud. *UCC*, (pp. 127-136).
- Madhani, S., Tauil, M., & Zhang, T. (2005). Collaborative sensing using uncontrolled mobile devices. *Int. Conf. on Collaborative Computing: Networking, Applications and Worksharing*, (p. 8).
- Mao, Y., Chen, Y., Hackmann, G., Chen, M., Lu, C., Kollef, M., & Bailey, T. (2011). Medical data mining for early deterioration warning in general hospital wards. *ICDMW*, (pp. 1042-1049).
- MAYO-PD. (2016). Parkinson's disease information (mayo clinic website). Retrieved from <http://www.mayoclinic.org/diseases-conditions/parkinsons-disease/basics/definition/con-20028488>
- Milenkovi, A., Otto, C., & Jovanov, E. (2006). Wireless sensor networks for personal health monitoring: Issues and an implementation. *Comput. Commun.*, 29(1314), 2521-2533.
- NVIDIA-DIGITS. (2016). NVIDIA DIGITS — Interactive Deep Learning GPU Training System. Retrieved from <https://developer.nvidia.com/digits>
- Olorode, O., & Nourani, M. (2014). Reducing leakage power in wearable medical devices using memory nap controller. *IEEE Dallas Circuits and Systems Conference (DCAS)*, (pp. 1-4).
- Page, A., Hassanaliheragh, M., Soyata, T., Aktas, M. K., Kantarci, B., & Andreescu, S. (2015b). Conceptualizing a Real-Time Remote Cardiac Health Monitoring System. In T. Soyata, Enabling Real-Time Mobile Cloud Computing through Emerging Technologies (pp. 1-34). IGI Global.
- Page, A., Kocabas, O., Soyata, T., Aktas, M., & Couderc, J. (2014). Cloud-Based Privacy-Preserving Remote ECG Monitoring and Surveillanc. *Annals of Noninvasive Electrocardiology (ANEC)*, 328-337.
- Page, A., Soyata, T., Couderc, J., & Aktas, M. K. (2015a). An Open Source ECG Clock Generator for Visualization of Long-Term Cardiac Monitoring Data. *IEEE Access*, 3, 2704-2714.
- Page, A., Soyata, T., Couderc, J., & Aktas, M. K. (2016). "QT clock" to improve detection of QT prolongation in Long QT Syndrome patients. *Heart Rhythm*, 13(1), 190-198.
- Page, A., Soyata, T., Couderc, J., Aktas, M. K., Kantarci, B., & Andreescu, S. (2015c). Visualization of Health Monitoring Data acquired from Distributed Sensors for Multiple Patients. *IEEE Global Telecommunications Conference (GLOBECOM)*. San Diego, CA.
- Pantelopoulos, A., & Bourbakis, N. (2010, Jan). A survey on wearable sensor-based systems for health monitoring and prognosis. *EEE Trans. Sys., Man, and Cybernetics, Part C: Applic. and Reviews*, 40(1), 1-12.
- Paradiso, R., Loriga, G., & Taccini, N. (2005, Sep). A wearable health care system based on knitted integrated sensors. *IEEE Trans. Info. Tech. in Biomedicine*, 9(3), 337-344.
- Park, C., Chou, P., Bai, Y., Matthews, R., & Hibbs, A. (2006). An ultra-wearable, wireless, low power ECG monitoring system. *IEEE BioCAS*, (pp. 241-244).
- Pedregosa, F. e. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825-2830.
- Pouryazdan, M., Kantarci, B., Soyata, T., & Song, H. (2016, Mar). Anchor-Assisted and Vote-based Trustworthiness Assurance in Smart City Crowdsensing. *IEEE Access*, 4, 529-541.
- Powers, N., & Soyata, T. (2015). AXaaS (Acceleration as a Service): Can the Telecom Service Provider Rent a Cloudlet ? *Proceedings of the 4th IEEE International Conference on Cloud Networking (CNET)*, (pp. 232-238). Niagara Falls, Canada.
- Powers, N., Alling, A., Osolinsky, K., Soyata, T., Zhu, M., Wang, H., . . . Kwon, M. (2015). The Cloudlet Accelerator: Bringing Mobile-Cloud Face Recognition into Real-Time. *Globecom Workshops (GC Wkshps)*. San Diego, CA.
- Rao, B. (2013). The role of medical data analytics in reducing health fraud and improving clinical and financial outcomes. *CBMS*, (pp. 3-3).
- Ray, P. (2014). Home Health Hub Internet of Things (H3IoT): An architectural framework for monitoring health of elderly people. *ICESMR*, (pp. 1-3).

- Rolim, C., Koch, F., Wastphall, C., Werner, J., Fracalossi, A., & Salvador, G. (2010). A cloud computing solution for patient's data collection in health care institutions. *ETELEMED*, (pp. 95-99).
- Sheng, X., Tang, J., Xiao, X., & Xue, G. (2013). Sensing as a service: Challenges, solutions and future directions. *IEEE Sensors Journal*, 13(10), 3733-3741.
- Siekkinen, M., Hiienkari, M., Nurminen, J., & Nieminen, J. (2012). How low energy is bluetooth low energy? comparative measurements with ZigBee 802.15.4. *WCNCW*, (pp. 232-237).
- SIGFOX. (2016). Global Cellular Connectivity for Internet of Things. Retrieved from <http://www.sigfox.com/en/>
- Son, D., Lee, J., Qiao, S., Ghaffari, R., Kim, J., Lee, J. E., . . . Kim, D. H. (2014). Multifunctional wearable devices for diagnosis and therapy of movement disorders. *Nature Nanotechnology*, 1-8.
- Soyata, T. (2015). Enabling Real-Time Mobile Cloud Computing through Emerging Technologies. IGI Global.
- Soyata, T., Ba, H., Heinzelman, W., Kwon, M., & Shi, J. (2013). Accelerating Mobile Cloud Computing: A Survey. In H. T. Kantarci (Ed.), *Communication Infrastructures for Cloud Computing* (pp. 175-197). Hershey, PA, USA: IGI Global. doi:10.4018/978-1-4666-4522-6.ch008
- Soyata, T., Copeland, L., & Heinzelman, W. (2016, Feb). RF Energy Harvesting for Embedded Systems: A Survey of Tradeoffs and Methodology. *IEEE Circuits and Systems Magazine*, 16(1), 22-57.
- Soyata, T., Muraleedharan, R., Ames, S., Langdon, J., Funai, C., Kwon, M., & Heinzelman, W. (2012). COMBAT: mobile-Cloud-based cOmpute/coMmunications infrastructure for BATtlefield applications. *Proceedings of SPIE*, (pp. 84030K-13). Baltimore, MD.
- Soyata, T., Muraleedharan, R., Funai, C., Kwon, M., & Heinzelman, W. (2012). Cloud-Vision: Real-Time Face Recognition Using a Mobile-Cloudlet-Cloud Acceleration Architecture. *Proceedings of the 17th IEEE Symposium on Computers and Communications (ISCC)*, (pp. 59-66). Cappadocia, Turkey.
- TensorFlow. (2016). Google TensorFlow. Retrieved from <http://www.tensorflow.org>
- Torfs, T., Leonov, V., Van Hoof, C., & Gyselinckx, B. (2006). Body-heat powered autonomous pulse oximeter. *5th IEEE Conference on Sensors*, (pp. 427-430).
- Ukis, V., Tirunellai Rajamani, S., Balachandran, B., & Friese, T. (2013). Architecture of cloud-based advanced medical image visualization solution. *CCEM*, (pp. 1-5).
- Weaver, C., Ball, M., Kim, G., & Kiel, J. (2015). *Healthcare Information Management Systems: Cases, Strategies, and Solutions*. Springer International Publishing.
- Wei, L., Kumar, N., Lolla, V., Keogh, E., Lonardi, S., Ratanamahatana, C., & Van Herle, H. (2005). A practical tool for visualizing and data mining medical time series. *Proc. 18th IEEE Symposium on Computer-Based Med. Sys.*, (pp. 341-346).
- WHO-COPD. (2016). World health organization factsheets: Chronic obstructive pulmonary disease (COPD). Retrieved from <http://www.who.int/mediacentre/factsheets/fs315/en/>
- WHO-CVD. (2016). World health organization factsheets: Cardiovascular diseases (CVDs). Retrieved from <http://www.who.int/mediacentre/factsheets/fs317/en/>
- Xu, S., Zhang, Y., Jia, L., Mathewson, K. E., Jang, K. I., Kim, J., . . . Rogers, J. A. (2014). Soft microfluidic assemblies of sensors, circuits, and radios for the skin. *Science*, 344, 70-74.
- Yang, D., Xue, G., Fang, G., & Tang, J. (2015). Incentive mechanisms for crowdsensing: Crowdsourcing with smartphones. *IEEE/ACM Transactions on Networking*, 1-13.
- Yu, C., & Sharma, G. (2010, Aug). Camera scheduling and energy allocation for lifetime maximization in user-centric visual sensor networks. *IEEE Transactions on Image Processing*, 19(8), 2042-2055.
- Yuriyama, M., & Kushida, T. (2010). Sensor-Cloud Infrastructure - Physical Sensor Management with Virtualized Sensors on Cloud Computing. *NBIS*, (pp. 1-8).
- Zhang, H., & Hou, J. (2005). Maintaining sensing coverage and connectivity in large sensor networks. *Ad Hoc & Sensor Wireless Networks*, 1(1-2), 89-123.
- Zhang, J., Iannucci, B., Hennessy, M., Gopal, K., Xiao, S., Kumar, S., . . . Rowe, A. (2013). Sensor Data as a Service - A Federated Platform for Mobile Datacentric Service Development and Sharing. *IEEE International Conference on Services Computing (SCC)*, (pp. 446-453).
- Zhang, Y., Meratnia, N., & Havinga, P. (2010). Outlier detection techniques for wireless sensor networks: A survey. *IEEE Communications Surveys and Tutorials*, 12(2), 159-170.
- Zhao, W., Wang, C., & Nakahira, Y. (2011). Medical application on internet of Things. *ICCTA*, (pp. 660-665).

Authors



Alex Page is a PhD candidate in Dr. Soyata's research group. He received his MS in Electrical Engineering from the University of Rochester. His current research is focused on computer systems for medical data analysis, including databasing, GPU acceleration, and machine learning techniques.



Burak Kantarci is an Assistant professor in the Department of ECE at Clarkson University, and the founding director of the next generation communications and computing networks (NEXTCON) research lab at Clarkson. He received his PhD in Computer Engineering at Istanbul Technical University in 2009. He is an editor for IEEE Communications Surveys and Tutorials. Dr. Kantarci also serves as the secretary of IEEE ComSoc Communication Systems Integration and Modeling Technical Committee (CSIM-TC). He is a senior member of IEEE and a member of ACM.



Shurouq Hijazi is an undergraduate student majoring in ECE with a concentration in Wireless Communications at the University of Rochester. She is currently a research assistant in Dr. Soyata's laboratory, and is focused on machine learning techniques.



Tolga Soyata received his B.S. degree in Electrical and Communications Engineering from Istanbul Technical University in 1988, M.S. degree in ECE from Johns Hopkins University in 1992 and PhD in ECE from University of Rochester in 2000. He joined the University of Rochester ECE Department in 2008. He was an Assistant Professor – Research at UR ECE until 2016. He joined SUNY Albany Department of CE as an Associate Professor in 2016. His teaching interests include CMOS VLSI ASIC Design, FPGA-based High Performance Data Processing System Design, and GPU Parallel Programming. His research interests include Cyber Physical Systems, Digital Health, and GPU-based high-performance computing. He is a senior member of both IEEE and ACM.



Dogan Askan is an MS student in the ECE department of Clarkson University. He is a researcher in Prof. Burak Kantarci's lab. His research interests include trustworthy sensing.