

Article

Making the Case for a P2P Personal Health Record

William Connor Horne and Zina Ben Miled * 

Department of Electrical and Computer Engineering, Indiana University Purdue University Indianapolis, Indianapolis, IN 46202, USA; wcfhorne@gmail.com or wchorne@iu.edu

* Correspondence: zmiled@iupui.edu

Received: 23 August 2020; Accepted: 24 October 2020; Published: 31 October 2020



Abstract: Improved health care services can benefit from a more seamless exchange of medical information between patients and health care providers. This exchange is especially important considering the increasing trends in mobility, comorbidity and outbreaks. However, current Electronic Health Records (EHR) tend to be institution-centric, often leaving the medical information of the patient fragmented and more importantly inaccessible to the patient for sharing with other health providers in a timely manner. Nearly a decade ago, several client–server models for personal health records (PHR) were proposed. The aim of these previous PHRs was to address data fragmentation issues. However, these models were not widely adopted by patients. This paper discusses the need for a new PHR model that can enhance the patient experience by making medical services more accessible. The aims of the proposed model are to (1) help patients maintain a complete lifelong health record, (2) facilitate timely communication and data sharing with health care providers from multiple institutions and (3) promote integration with advanced third-party services (e.g., risk prediction for chronic diseases) that require access to the patient’s health data. The proposed model is based on a Peer-to-Peer (P2P) network as opposed to the client–server architecture of the previous PHR models. This architecture consists of a central index server that manages the network and acts as a mediator, a peer client for patients and providers that allows them to manage health records and connect to the network, and a service client that enables third-party providers to offer services to the patients. This distributed architecture is essential since it promotes ownership of the health record by the patient instead of the health care institution. Moreover, it allows the patient to subscribe to an extended range of personalized e-health services.

Keywords: personal health record; peer-to-peer; service oriented architecture; Bayesian networks

1. Introduction

Access to the complete medical history of the patient is becoming increasingly crucial for the delivery of quality medical services. In order to provide this access, extensive efforts have been devoted to the digitization of medical records [1] and to facilitating the exchange of these records among health care providers. In the United States, the Centers for Medicare and Medicaid Services (CMS) and the Office for the National Coordinator for Health Information Technology (ONC) actively promote the adoption of Electronic Health Records (EHRs) through Medicare and Medicaid programs [2,3]. Example commercial EHR systems that have been deployed in various health care institutions include EPIC [4], Cerner [5], and Meditech [6]. These systems are able to digitize processes and workflows. However, they each use different data representations which makes it difficult to seamlessly exchange health records between institutions that use different EHRs. In order to overcome this interoperability gap, standard data representations and ontologies (e.g., HL7 [7], FHIR [8]) have been developed. In addition, Health Information Exchanges (HIEs) [9] emerged as third party brokers that aggregate health records from various providers, translate these records to a standard data representation, and facilitate their delivery to requesting parties [10].

As of 2017, 69% of hospitals in the United States had agreements with an HIE [11]. However, this number can be misleading. The coverage of HIEs is often regionally limited. As a result, providers need to subscribe to multiple HIEs. Similarly, in Europe, countries autonomously manage their individual health care systems and have the same interoperability concerns. In order to address this challenge, the ONC sponsored the eHealth Exchange [12] for interstate exchanges in the US and the European Union is targeting the seamless transfer of prescription and patient summaries among its member countries by 2021 [13] through the eHealth Digital Service Infrastructure (eHSDI) [14]. While these efforts are expected to substantially enhance the access to more complete health records, they will not be able to provide full coverage for all patients. For instance, HIEs may not include small health providers [15]. Moreover, EHRs and HIEs may not be able to accurately aggregate records from multiple sources because of the lack of a unified patient identifier. In fact, HIEs often rely on heuristics in order to match records with error rates reaching up to 38% [16].

Essentially, HIEs do not fully support the ability of a patient to maintain a complete health record and to freely share this record with any health care provider. They are institution-centric rather than patient-centric. For instance, patients may want to transition from one provider to another due to travel or relocation extending beyond the boundaries of a single health network, state or country. Comorbidity may necessitate the interaction with multiple health service providers. Moreover, emergency health events may require quick access to patient health care information. Finally, patients may want to subscribe to e-health services, such as personalized risk prediction services for chronic diseases. This type of flexibility can only be achieved with a patient-centric health information system that allows patients to manage their own health records.

Previous efforts at providing patient-centric health information systems include portals and PHRs. Portals [17] suffer from the same limitations mentioned above for EHRs because they are tightly coupled with the health institution that manages the portal [18,19]. Several PHRs solutions were proposed nearly a decade ago [20]. These PHRs faced resistance from different stakeholders, uptake was limited and most have since been discontinued. Some of the challenges associated with these previous PHRs include:

- Implementation and deployment difficulties [21–23];
- Concerns from patients about the security of their health records [21–23], and the ability to adequately manage them [22,24];
- Concerns from physicians related to the interpretation of the health information by the patient and the validity of the record being held by the patient [21,23,24];
- Ownership of the health data by the patient and the impact this may have on patient retention for health care institutions [21,25–27].

A decade later, the increasing demand for more flexible health services [28], the improved health literacy [29] of the patients and recent technology advances indicate that the potential benefits of a new PHR system may outweigh the challenges that previous PHRs faced. However, the new PHR must take into consideration the lessons learned and the pitfalls of previous institution-centric and patient-centric health information systems. This paper proposes a new PHR architecture and demonstrates the ability of this architecture to fulfil the functional needs of the patients, the health care institutions and third-party e-health services.

For instance, previous PHRs adopted a client–server architecture. Moreover, they did not offer functionalities for flexible patient–providers data exchange. This design decision made the patient responsible for the aggregation of health data from different sources. On the other hand, HIEs focused on inter-institution data exchange. This led to gaps in the health records of the patients. More importantly, PHR, EHR and HIEs alike did not plan for the increasing demand from patients for personalized e-health services as well as the proliferation of home health devices. Recently, several researchers have been focusing on developing risk predictors for chronic diseases (e.g., dementia [30]). Others are proposing health monitoring frameworks for chronic health conditions,

such as diabetes [31,32]. These frameworks leverage home health devices and social media posts in addition to EHR data. Most of these emergent models rely on machine learning techniques and show promise for personalized preventive health services. However, they do not provide a clear pathway on how the services can be delivered to the patient. Specifically, who is responsible for obtaining and maintaining the data needed by the patient to invoke the services?

In the example of the risk prediction for dementia, the machine learning model requires the medical history of the patient. This includes prescription types and frequency, medical notes and diagnosis over a period of ten years. This information is only available from multiple EHRs or HIEs, especially if the patient suffers from multiple chronic diseases (e.g., diabetes and arrhythmia). EHRs records for each condition may reside with different EHRs and potentially different HIEs. Moreover, the patient may have moved during the ten year period making his or her health record split across multiple health institutions. This record fragmentation is more prevalent among the older adult population [30]. Unfortunately, this is the patient population that needs this service the most.

The above example highlights the longitudinal and multi-institution data aggregation problem when only one type of data (i.e., EHR data) is required to invoke the e-health service. The problem is exacerbated when the prediction model requires data from additional sources. For example, the diabetes health monitoring framework [31,32] mentioned above relies on social media posts from the patient and home health monitoring device data in addition to EHR data. The authors offer a federated framework that can fuse the data from these three sources and monitor the health status of diabetic patients. Deploying such a system in the community entails that all health providers, patients and third party services subscribe to the same framework. This is one of the pitfalls of the federated approach that was adopted by HIEs. Alternatively, the service provider may establish a partnership agreement directly with a single health institution. In this case, the patient will be restricted to the services offered through the health institution. Moreover, the third party service may not have access to the entire health record of the patient, which may be spread across multiple institutions.

This paper proposes a distributed PHR platform that overcomes the data accessibility and sharing issues of current systems and enables patients to customize health services to their own needs. Indeed, this platform allows the patients to:

- Maintain a complete health record;
- Complement this record with data from other sources (e.g., home health devices or social media);
- Share this data with the health provider of their choice;
- Subscribe to the e-health services that match their health conditions.

The remainder of the paper is organized as follows. Section 2 provides an overview of previous related work. Section 3 describes the proposed P2P PHR. Section 4 presents examples of the workflows that are supported by the proposed system. Section 5 concludes by summarizing the main advantages of the proposed platform as well as directions for future work.

2. Background

As discussed above, there is a strong motivation for a PHR system. In fact, a PHR is a key enabler of e-health and a core component of any future health network redesign [33]. In this section we review previous related work with respect to the architecture design choices and semantic interoperability of health information systems. We also highlight the specific issues that the proposed model attempts to address with respect to these two areas.

The successful deployment of a PHR depends on the implementation strategy which is primarily guided by the underlying architecture of the system. There are two main potential architectures; namely, client-server and peer-to-peer (P2P). In a client-server architecture, clients connect to a central server [34] and exchange records directly with the server. In a P2P network, clients are called peers and communicate directly with each other [34]. Each peer must operate as both a client and a server in the network. Most of the early PHR solutions were based on a client-server architecture. Some of

these solutions are described in [20,35,36]. Only few recent health information systems use a P2P architecture. For instance, a P2P framework using a tuple space architecture [37] to exchange records across HIEs was proposed in [38]. A P2P PHR with subnetworks was proposed in [39]. A PHR focusing on medication history using blockchain over a P2P network was introduced in [40].

P2P networks follow a distributed architecture and have several advantages over client–server architectures. They are self-scalable [34], allowing growth to hundreds of thousands of peers. Compute, network, and storage cost are shared by the peers and not concentrated in a single server. P2P networks are also resilient against network failures [41]. For the proposed PHR application, these characteristics are necessary because they support the efficient scaling of the system to a large number of patients and promote the distributed ownership of the network.

Compared to the client–server architecture, the main limitations of P2P networks are mostly related to implementation difficulties. These include:

- Connection issues due to symmetric network address translation (NAT) [42]. Symmetric NAT is a routing method that changes the source and destination IPs;
- Peer-churn [34], where peers enter and leave the network too quickly, leading to network instability.

In addition, transactions in P2P networks are significantly different from the traditional OLTP transactions used in client–server architectures. Transactions in P2P networks are based on consensus. Solving the consensus problem in a bounded time is difficult. However, solutions such as the Two-Phase-Commit [43], Paxos [44], Byzantine consensus [45] and Blockchain [46] have recently been proposed to adequately address this issue. In this paper, a modified Two-Phase-Commit protocol is used to address the consensus problem in the proposed P2P PHR network. Adequate solutions for the peer-churn and connection issues are discussed in [47].

From a semantic perspective, the exchange of health data among different stakeholders, whether in a federated system as in the case of HIEs or a distributed system as in the case of P2P platforms, requires the participating systems to use common data representation. Semantic interoperability among heterogeneous information systems is a long standing research area in various sectors including the health sector. The research community has mostly converged towards the used of common ontologies and standards in order to facilitate this exchange.

For example, in [48], the authors use ontologies to map medical concepts in a source EHR system to a target EHR system. This approach avoids complex transformation procedures that may not be sustainable since the source or the target EHRs may change their representations or encoding of the medical records over time. In this application, the semantic transformation occurs at the target system. Alternatively, semantic mapping can be performed at the source as demonstrated in [49]. In this latter case, the health record is mapped to a common ontology at the source prior to the data exchange.

Ontologies were also used in a decision support system [50] to enable the sharing of patient's health condition. This decision support application was within the context of ambient assisted living and demonstrates the emergent need for data exchange among clinical and non-clinical stakeholders [50,51]. While focusing primarily on semantic interoperability, this application also underscores the need for information flow between smart homes and the more traditional clinical or medical environments. Similarly, in [52], the authors discuss the difficulties associated with caring for older adults and the multiple stakeholders that participate in their care plan. A personalized ontology that is specific to each patient is proposed. This ontology helps the health-care professionals to cooperate towards an integrated health care plan for the patient.

The applications discussed above help illustrate the complexity of current medical services and the importance of a coordinated health plan. This plan may only be achievable through a personalized health platform. The above examples focus on semantic interoperability. Our proposed PHR focuses on the workflows underlying the exchange of the health data and on enabling access to a wide range of health services. Our proposed system uses the most recently accepted ontologies and standards; namely the Fast Healthcare Interoperability Resources (FHIR) [53] and ICD-10 [54]. However, future work

may consider the convergence of these two endeavors by utilizing personalized ontologies over the proposed P2P PHR. We anticipate the personalized ontologies to be especially usefully when health data are extended to include, for example, social media posts from the patient.

3. Methods

The proposed PHR is a distributed system that allows patients to communicate with different health providers and subscribe to various e-health services. By using the proposed platform, the patients can manage and share their records with the stakeholders involved in their health care plan. This patient-centric view of health care management requires a platform where the ownership of the health data resides with the patient. Moreover, the platform must allow for a large number of participants and transactions. This section describes the design choices made during the development of the proposed system in order to achieve the aforementioned goals.

As discussed in Section 1, previous client–server PHR and federated HIEs have limitations in terms of longitudinal coverage, scalability and access flexibility. Therefore a distributed P2P architecture was selected. P2P networks fall under three categories: *centralized*, *pure*, and *hybrid*. Network design and topology play a key role in the network performance. In fact, network scalability and management difficulties increase from *centralized*, to *hybrid*, to *pure*.

The high level network architecture of the proposed system is shown in Figure 1. It follows the *centralized* P2P design and consists of: an index server, a peer client, and a service client. The choice of a centralized P2P network over hybrid or pure P2P network architectures was primarily motivated by:

- The ease of deployment of the architecture;
- Its ability to enforce privacy measures according to HIPAA or other health regulations through the index server;
- The fact that it represents a path of least resistance to change from the current institution-centric EHRs or HIEs.




Figure 1. System Architecture.

The index server provides resource lookup for the network and mediates network access [41]. All network participants are required to register in the index server. This information is stored in the registration database on the index server and is indexed by the unique ID of the participant. The index server uses this information for authentication purposes. The index server also maintains the history and current status of all transactions in a different database as shown in Figure 1. This allows the index server to maintain the protected health information of the patient (e.g., name, address, phone number) separate from other operational data. The separation is important since protected health information

must adhere to strict privacy rules (e.g., HIPAA). We should also note that the second database in the index server (Figure 1) only includes the status and metadata of the transactions. This database neither stores medical records nor the participants' protected health information. The former is directly exchanged between peers without the involvement of the index server. The latter is stored in a separate database and only exchanged between the participant and the index server during network registration.

There are also two types of clients in the network: peer clients and service clients (Figure 1). Peer clients are used by health providers and patients to manage records, upload document metadata to the index server, search for network resources, request transactions, and connect to other peers. The service client is used by third-party health services to offer services to the participating patients in the network. As mentioned in Section 1, a wide range of services are available and more can be anticipated in the future. Example e-health services include predicting patients' risk for chronic diseases [30] or medication non-adherence [55]. Other e-health services may focus on monitoring the health condition of chronically ill patients [50–52].

The following subsections describe the components of the proposed system. The representation used for the patient health record is discussed first. The following subsection describes the transactions that enable the exchange of information between the participants. Finally, the third subsection demonstrates the ability of the platform to interact with third-party e-health services.

3.1. Records

A patient's electronic health record documents the health of a patient and his or her interactions with healthcare providers. In the proposed system, these records are codified in the FHIR standard [8,56]. FHIR was recently proposed as an exchange standard that can facilitate the transfer of records between providers. The standard has two major sections: (1) a generalized ontology that represents the various information in an EHR and (2) a Representational State Transfer (REST) protocol that allows for the exchange of this information. The standard was previously used to develop several applications, including EHR systems [57,58].

The ontology component of the FHIR standard organizes the information captured in a health record. It details roles, workflows, and financial information as well as observations, conditions, and encounters. Condition diagnoses are often standardized according to the International Classification of Diseases (ICD-10 [54]). In summary, the entities in an FHIR ontology [53] include:

- Individuals and their roles in the system (e.g., patient, provider, etc.);
- Organizations, locations or devices;
- Workflows (e.g., tasks or appointments);
- Encounters between a patient and a health care provider;
- Clinical information such as observations, conditions, and medications;
- Financial information including billing.

In the proposed system, the patients and providers store their health records locally on their own computer. Each record is associated with a metadata, which consists of the origin, the creation time, and a SHA256 [59] hash of the record that establishes its uniqueness. Only the metadata are uploaded to the P2P network for discovery. Access to the actual record is granted through an established transaction protocol and after the approval by the data owner, as described next.

3.2. Transactions

The proposed system is transactional. Patients selectively authorize access to their records which are stored in the FHIR format. The index server logs each transaction, thereby creating an access trace. Because of the distributed architecture of the proposed framework, achieving consensus among participants when executing transactions is difficult. This consensus is easy to implement in a client-server model since all transactions are between two parties where one party is always

the server and this latter one-sidedly determines the status of the transactions. In order to achieve consensus in the proposed distributed architecture, a modified version of the Two-Phase-Commit protocol [60] is used. Under this modified protocol, the index server manages the global state of the system and facilitates synchronization among clients. Upon connecting to the network, clients replicate the state of the index server. Subsequently, they can post updates to the server. The server checks these updates against its own internal copy of the global state and periodically broadcasts state updates to all clients. This protocol is typically used in distributed simulations [60,61]. The choice of this modified Two-Phase-Commit protocol was motivated by the benefits it offers with respect to security, correctness, and resilience. These aspects are essential for the proposed PHR application. In fact, under this protocol there are no assumptions about an implicit trust of the clients, incorrect actions are not possible, and client failures do not block the progress of the server. These benefits come at the expense of acceptable limitations. For instance, clients may have to repeat actions if they are out-of-date with respect to the global state of the index server.

The network transactions are modeled as a state machine and the state of each transaction changes as predefined conditions are met. The index server acts as the ground truth between two communicating peers and ensures that the transactions are both valid and consistent. In addition, the index server provides an independent log of the transactions for audit purposes.

Three operations are supported in the proposed system: *Transfer Request*, *Push*, and *Service*. The first two are implemented using REST operations. The third was added to the REST protocol with minimum deviation from the standard in order to enable third-party services. The *Transfer request* operation (Figure 2) allows a client (source client) to request a health record from its owner (target client). This operation is used to exchange health records between a patient and his or health provider. The source client can be the patient and the target can be the health provider. For instance, the patient may have recently undergone a test or a procedure and is requesting the related report. Similarly, a health provider may request the health record of a patient that was created by another health provider. For example, a cardiologist may request the medical report that was obtained by the patient for his or her diabetes condition.




Figure 2. Transfer Request Operation. This operation is used by the patient to request a health record from the health provider or by the health provider to request a health record from the patient.

The implementation of the transfer operation follows a typical *get* REST operation (Figure 2). In order to maintain the global status of the network up-to-date, both the source and the target clients must send an update to the index server. The *push* operation, shown in Figure 3, allows a client to send a health record to another client and is implemented following the specifications of the *push* REST operation. For instance, patients may want to send an update to one of the health providers involved in their care management plan. Similarly, a health provider may want to update a patient on the result

of a laboratory test. The difference between a transfer request operation and a push operation is that the former is initiated by the requesting party, while the latter is unsolicited. In both cases, the two clients involved in the transaction have to send an update to the index server. However, in the case of the push operation, the initiating client sends the health record to the target client and the target client responds with an acknowledgment, as shown in Figure 3.




Figure 3. Push Operation. This operation is used by either the patient or the health provider to update the health record maintained by the other party.

The *Service* operation is initiated by the patient to request a service from a third party. For instance, the patient may want to know his or her specific risks for diabetes. The patient is able to select one or more service from those registered in the network. These services have to publish the description of the required health data needed to invoke their service. The service operation is then used by the patient to send this data to the selected service peer. The service provider performs the requested service and returns the results to the patient, as shown in Figure 4. As in the case of the previous transactions, both the requesting client and the service client must send an update to the index server. In addition, the exchange of health data is restricted between the patient and the service provider.




Figure 4. Service Operation. This operation is used by the patient to request an e-health service from a third party.

The service operation is needed to deliver e-health services to the patient. In fact, most current literature presents a conceptualization of these services in isolation from other services [30,55] or imposes the consolidation of all the health data related to each specific patient using a client–server

approach [50,51]. These approaches either do not provide a clear workflow for the patient to access the proposed e-health service or assume a client–server model that may fall into the same limitations as those of current HIEs. In the next section, we describe the steps that are involved in the development of an example e-health service and its deployment within the proposed network. The selected e-health service is hypertension risk prediction. However, the procedure is applicable to other e-health services.

3.3. Hypertension Predictor

In order to demonstrate the potential deployment of third-party services, a hypertension risk prediction service was implemented within the proposed platform. This service is not intended to showcase state-of-the-art machine learning models. There are several research studies that explore advanced machine learning models for diagnosis and risk prediction of chronic diseases [30–32]. The hypertension risk prediction service discussed in this section is used to illustrate the steps needed to develop an e-health service and to integrate it with the proposed platform. This section also discusses the gap that prevents currently proposed e-health services from being accessible by all the patients. Filling this gap is one of the objectives of the proposed framework.

Hypertension risk prediction was selected because this chronic disease has a high incidence rate and is more likely to be one of the earliest services to be offered. There are typically two phases in the development of a machine learning model. The first phase is the training of the model and the second phase is the validation of the model. These two phases are offline procedures. Once the training is completed, the second phase evaluates the accuracy of the model. As previously mentioned, typically previous research studies either do not consider deployment procedures or assume that the model will be deployed in a client–server model (e.g., within an HIE). In the remainder of this section, the training and testing phases of the hypertension service are first presented. A procedure that makes this service available to each patient is then discussed.

3.3.1. Model Development

The hypertension prototype service utilizes a Naive Bayes [62] model. It is trained with data extracted from the Medical Expenditure Panel Survey (MEPS) [63]. Naive Bayes machine learning models are attractive because they are interpretable [64] and can be bootstrapped with knowledge from domain experts [62]. However, the procedure presented in this section can accommodate any machine learning model, including those built on novel machine learning techniques [32]. The Naive Bayes model [65] assumes that each input feature is conditionally independent from other features. The risk for hypertension is expressed as the product of the conditional probabilities of the input features, as shown in (1).

$$P(C, E_n) = P(C) \prod_i P(E_i|C) \quad (1)$$

where C is the risk for hypertension and E_n represents the model input features.

The Medical Expenditure Panel Survey (MEPS) [63] is a yearly survey of households by the Agency for Healthcare Research and Quality (ARHQ). ARHQ's goal is to survey current patient healthcare in the United States. Each survey consists of questions about conditions, providers, costs and medication. Moreover, each cohort of households is followed over a period of two-years called a *panel* with five rounds in each panel. Twenty panels are available in the MEPS dataset starting from 1996. Each panel includes approximately 15,000 households. However, only a subset of the participants in the panel's first year are carried over to the second year. Moreover, as in the case of other datasets, there are several missing data values in the MEPS dataset.

As expected, from all the disease conditions reported in the MEPS dataset, hypertension had the highest record count (e.g., 3078 instances in Panel 19). This data availability and accessibility to heart-related information, motivated the choice of hypertension risk prediction as the e-health service pilot. A review of the literature helped prune the list of available pre-existing conditions in the MEPS

dataset to those that are relevant to hypertension risk prediction. The ICD-9-CM codes of 30 heart and respiratory conditions were selected, as shown in Table 1. It should be noted that, in order to protect the privacy of the survey participants, ARHQ only reports the first three characters of the ICD-9-CM code [66].

Table 1. Selected Pre-existing Conditions.

ICD-9-CM Group	Meps Code	Feature
410-414: Heart Disease	410	Acute Myocardial Infraction
	413	Angina Pectoris
	414	Other forms of Ischemic Heart Disease
415-417: Heart Failure	415	Acute Pulmonary Heart Disease
420-429: Circulatory System Diseases	424	Other Diseases of Endocardium
	425	Cardiomyopathy
	427	Cardiac Dysrhythmis
	428	Heart Failure
	429	Ill-defined Descriptions and Complications of Heart Disease
430-438: Cardio Brain Hemorrhage, Stroke	436	Acute, but ill-defined, Cerebrovascular Disease
440-449: Restricted Blood Flow	440	Atherosclerosis
	441	Aortic Aneurysm and Dissection
	442	Other Aneurysm
	443	Other Peripheral Vascular Disease
	444	Arterial Embolism and Thrombosis
	447	Other Disorders of Arteries and Arterioles
451-459: Issues with Veins and Lymph System	454	Varicose Veins of Lower Extremities
	455	Hemorrhoids
	458	Hypotension
	459	Other Disorders of Circulatory System
490-496: Chronic Respiratory Issues	490	Bronchitis, not specified as Acute or Chronic
	491	Chronic Bronchitis
	492	Emphysema
	493	Asthma
	496	Chronic Airway Obstruction, not elsewhere classified
510-519: Secondary Respiratory Conditions	511	Pleurisy
	514	Pulmonary Congestion and Hypostasis
	518	Other Diseases of Lung
	519	Other Diseases of Respiratory System

The above pre-existing conditions were augmented with additional health and demographic information. MEPS has only a few direct measures of the health status of the patient (e.g., Body Mass Index). However, the survey includes questions related to advice from doctors, major health events (e.g., stroke), and smoking habits [66]. The answers to these questions were used to extract an additional 17 features as shown in Table 2. The dataset used to train and test the hypertension risk prediction model is comprised of the combination of the features in Tables 1 and 2.

The development of the hypertension prediction model is performed in two steps: training and testing. This study used panels 17, 18, and 19, which were conducted from 2012 to 2015. These panels were selected for their consistency as other panels have significant variances in their survey questions. As mentioned above, each panel spans two years and is conducted in five rounds. These rounds were used to create the prior and post conditions needed to train and test the model. Indeed, the first half of each panel (i.e., rounds 1, 2, and part of 3) is used to create the input features of the model. The second

half of the panel (i.e., the remaining part of round 3 and rounds 4 and 5) are used to establish the outcome for each patient (i.e., whether or not the patient has hypertension).

Table 2. Selected Demographics and Health Condition Features.

MEPS Feature	Description
SEX	Sex
RACE1VX	Race
AGELAST	Person's age last time eligible
HIBPDX	High Blood pressure diagnosis
CHHDX	Coronary Heart Disease diagnosis
ANGIDX	Angina diagnosis
MIDX	Heart Attack (MI) Diagnosis
OHRTDX	Other Heart Disease Diagnosis
STRKDX	Stroke Diagnosis
EMPHDX	Emphysema Diagnosis
CHOLDX	High Cholesterol Diagnosis
DIABDX	Diabetes Diagnosis
ADSMOKE42	SAQ: Currently Smoke
NOFAT53	Restrict High Fat/Cholesterol
EXRCIS53	Advised to exercise more
BMINDX53	Adult Body Mass Index
POVCAT	Family Income as Percentage of Poverty Line

A data pipeline was used to convert the MEPS dataset into features and outcome according to the above mentioned split for each patient. The data were also imputed. Missing disease conditions or ICD-9-CM code from the input features were set to 0. We realize that this may lead to under reporting. However, it was unavoidable as the MEPS dataset does not have definite no answers in many cases. In addition, numerical input features such as body mass index (BMI) and age are binned to categorical classes. The resulting dataset had a total of 26,653 patient records. It was split into training and testing datasets using a 70/30 split. It should be noted that the hypertension prediction model, as described in this paper, is for illustration purposes. In order for the model to be useful in production it should be trained on a higher quality data and with a larger number of patients.

3.3.2. Model Deployment

In order to invoke the above hypertension risk prediction model for their specific case, the patients need to supply the required values for the input features of the model. There are eight features related to pre-existing condition (Table 1) and 17 features related to demographics and other medical conditions (Table 2). When the proposed PHR is used, this information is available to the patient and can be provided to the hypertension service using the service transaction described in Figure 4. Moreover, patients can subscribe to any e-health service they choose. This competitive market place will help improve the e-health service offering. For instance, the hypertension risk prediction service described above will have to publish the dataset it was trained with and the accuracy obtained on the test dataset. These metrics can be used by the patients to compare and contrast different services.

4. Implementation and Results

The proposed framework was developed and deployed in a laboratory environment. This section describes the details of this implementation and illustrates the functionalities of the system with examples of transfer, request and service transactions. The index server, peer client, and service client are implemented using the Go programming language [67]. The data are stored in MongoDB [68]. MongoDB was selected over an SQL database because it can support collections of documents. The framework also uses the MongoDB driver (mgo [69]) and the FHIR intervention engine [70] libraries to connect the application layer with the backend database.

The index server manages the network on behalf of the peers. It performs the signaling, signup, and lookup services in the network. It also tracks the registered users, peer IP addresses, and the status of the transactions and services. The peer client is a local application with a Web-User-Interface that allows patients and providers to interact with the network and manage health records. Using this application, users can register in the network, login, lookup resources, issue requests for transactions, and post the metadata of health records to the network. All of these operations are implemented as REST [71] APIs. The service client is used by the service providers. This client periodically checks the network for new requests for its service. Because of the sensitivity of health information, the service client performs the requested service but does not retain the input or output data. This restriction can be enforced when the service is initially registered in the network.

4.1. Client Transactions

There are two main client transactions: transfer request (Figure 2) and push operation (Figure 3). All the operations in the framework use a design template with the same structure. Each operation is a state machine with the following states: *waiting*, *pending*, *final*, *failed*, and *anceled*. We illustrate the implementation of the transfer request protocol. The push operation uses a similar protocol.

In the case of the transfer request, the *waiting* state indicates that a transfer request has been initiated but the owner of the document has not yet processed the transaction. The *pending* state indicates that the owner approved the transaction but processing of the transaction is not complete. The *final* state indicates that the transaction completed successfully. The *failed* state corresponds to a failed transaction. The *anceled* state informs the client that issued the request that the transaction was canceled. While this structure is the same for all the operations that are performed by the clients, the transition from one state to the next will vary depending on the selected operation.

Figure 5 shows the UI that is used by the client to initiate a transaction or invoke a service. In order to initiate a transfer request, the peer client must send a request to the index server. The request includes the document hash and the target peer. These are the first two fields under the record metadata in Figure 5. Once a request is issued, the server validates the request and creates a new transaction, which is placed in the *waiting* state. The target peer will then review the list of *waiting* transactions (Figure 6) and decide how to process these transactions. It can approve, reject, or ignore each request. The example in Figure 6 shows that the request is waiting for the authorization of the target peer. The remaining fields in the figure are the unique id of the transaction in the network, the hash value of the requested document, the id of the requesting peer, the id of the target peer, the transaction type and the time stamp of the transaction.

The decision of the target peer is sent to the index server via a post method. If the target peer rejects the request, the server changes the state of the transaction to *anceled* and no further action is required. If the request is approved, the index server moves the transaction to the *pending* state. Both the requesting and the target peers can then pull the transaction from the server and add it to their respective job lists. All the peers continuously process jobs from their job lists and periodically poll the server for transaction updates. In order to process the pending transaction, the requesting client establishes a TCP connection with the target peer. The target peer uses this connection to transfer the requested document directly to the peer that initiated the connection. Once the document transfer is completed, both peers inform the index server and the state of the transaction is changed to final as shown in Figure 7.

The index server also operates as a state machine. It continuously checks the incoming messages from the peers and evaluates whether or not the conditions for a state progress of a given transaction are satisfied. In the above example, the index server moves the transaction from the waiting (Figure 6), pending to the complete state (Figure 7). If there is no progress on a transaction, a failure occurs, and the clients will need to retry the transaction.

Transaction Request

Record Metadata

```
{5d356ae675c13beebb53acc2 5d35692175c13beebb53a5ab 8a7dfd023ce520c27b5fcb8bff55dcdf1aad1562db8f161ab31e0a846a6422cc 2019-07-22 07:50:39.868 +0000 UTC local}
```

Basic Transactions

Choose an option

Transfer

Push

Submit

Services List

Service List

Id	Name	Uid	Heartbeat	Info	Data Type
<input checked="" type="radio"/> 5cbf35247cd67aeebff7ea68	meps	5cbf35247cd67aeebff7ea65	2019-07-22 07:51:35.231 +0000 UTC	info about the meps predictor	heartinfo

Submit

Figure 5. The UI used by the patient to initiate a transaction or invoke a service.

Id	DocHash	SenderId	ReceiverId	Action State	Time	Authorize
5d356a1844e6be19d2a9a5c6 a54d1503a72385cdb74e9ca93cce4ef26128e70d35074f02b1dca20e5bc4cf9f 5d3569d375c13beebb53a84b 5d35692175c13beebb53a5ab				transfer waiting	2019-07-22 07:47:36.465 +0000 UTC	Authorize Transaction

Figure 6. List of pending transactions in the network. This list is maintained by the index server.

Id	DocHash	SenderId	ReceiverId	Action State	Time	Authorize
5d356a1844e6be19d2a9a5c6 a54d1503a72385cdb74e9ca93cce4ef26128e70d35074f02b1dca20e5bc4cf9f 5d3569d375c13beebb53a84b 5d35692175c13beebb53a5ab				transfer final	2019-07-22 07:47:36.465 +0000 UTC	

Figure 7. List of completed Transactions. The status of the transactions is maintained and updated by the index server.

The record that is exchanged between the requesting and the target peers is in json FHIR format. An example record is shown in Figure 8. It consists of two parts—the metadata and the actual content of the document. The first fields in the metadata are the same as in Figures 6 and 7. The last two fields of the metadata are the time stamp of the transaction and the source of the document. This record is exchanged directly between the two peers involved in the transaction. It is not accessible to the index server. The index server only has access to the metadata.

Record:
a54d1503a72385cdb74e9ca93cce4ef26128e70d35074f02b1dca20e5bc

Metadata

```
5d35692175c13beebb53a5ab 5d356a1844e6be19d2a9a5c6 a54d1503a72385cdb74e9ca93cce4ef26128e70d35074f02b1dca20e5bc4cf9f 2019-07-22 07:44:59.403 +0000 UTC http://hapi.fhir.org/baseDstu3
```

Data

```
{ "id": "914869", "birthDate": { "precision": "date", "time": "2018-12-16T05:00:00Z" }, "gender": "male", "meta": { "lastUpdated": { "precision": "timestamp", "time": "2018-12-16T19:41:36.807Z" }, "versionId": "1" }, "name": [ { "family": "Walker", "given": [ "Austin" ], "text": "Austin Walker", "use": "official" } ], "resourceType": "Patient", "text": { "div": "\u003cdiv xmlns=\u201chttp://www.w3.org/1999/xhtml\u201c\u003eAustin Walker\u003c/div\u003e", "status": "generated" } }
```

Figure 8. Example data record. The first field is the unique record id; the second field is the record metadata; the third field is the content of the record.

4.2. Service Transactions

The service transactions are different from the peer transactions. The bottom of Figure 5 shows the list of services that are registered in the network. These services are available to the patient. Each service is identified by a service id, a name, and the id of the service provider. The list of services

also includes the timestamp of the last time the service was pinged by the index server, human readable information about the service and the type of data the service expects.

Once a patient issues a service request through his or her client, the input data expected by the service are transferred to the service peer and the corresponding service is invoked. In the case of the hypertension risk prediction service, the input data are the values of the features of the hypertension prediction model. The result generated by the service is returned to the requesting peer as shown in Figure 9. In this case, neither the metadata nor the content of the document are available in the network. The patient can elect to post the metadata of the document to the network as shown in the bottom of Figure 9. The last id in the metadata of the document is the id of the owner; in this case the hypertension service provider.

Local Record:
8dc8f3631ae6036ec5b777070e72545c01c83267ecec7a93b8f0473

Metadata
8dc8f3631ae6036ec5b777070e72545c01c83267ecec7a93b8f04736a44a4885 2019-07-22 07:52:46.18 +0000 UTC 5d356b4444e6be19d2a9a5c7

Data
{ "result": "true" }

Network Options

Figure 9. Result record return by the hypertension service in response to the patient's request.

The hypertension risk prediction model was deployed as a service in the network. The accuracy, precision, sensitivity and specificity of the model are included in Table 3. Despite the fact that the model was trained on small dataset, Table 3 shows that the accuracy of the model is acceptable. When deployed in production, this service must be trained on a large dataset and its accuracy, precision, sensitivity and specificity should also be published as part of the service metadata. This allows the patients to make an informed evaluation of the quality of the service.

Table 3. Performance of the Hypertension Risk Prediction Model.

Performance Metric	Percentage (%)
Accuracy	75.2
Precision	58.9
Sensitivity	66.3
Specificity	79.1

The above protocol was demonstrated for a single service. However, the protocol can be extended to other e-health services. The intent of the protocol is to allow the patients to subscribe to a customized list of services that are aligned with their needs and medical conditions.

5. Conclusions

The framework introduced in this paper is a prototype of a P2P personal health record network. It can help each patient maintain a complete health record, share specific documents from this record with health providers, and easily subscribe to third party services. We believe this framework can help make the delivery of health services more efficient and more accessible to the patients. Moreover, we believe that the framework can also help make personalized predictive medicine available to all. While there are no guarantees for wide adoption, the proposed PHR improves the potential for adoption by patients since it provides them with more flexibility. Moreover, during the

past decade the patient literacy, digital ability and demands for flexible health care services have increased considerably, providing an opportunity for the re-launch of PHR systems.

The proposed framework is supported by a REST protocol over a distributed P2P network. The network is managed by a central index server. Peers interact with the network by using two types of transactions. The first type facilitates the exchange of documents between peers and the second allows a peer to request a service from a service provider. In order to allow resource discovery, peers must post the metadata of their documents to the network and services must register in the network. The actual exchange of health records among peers and with service providers occurs directly between two peers without the involvement of the index server. This design choice is necessary to ensure that the only copies of the health records that are persistent are those authorized by the owner.

The prototype was developed in a laboratory environment as a proof of concept. Large scale testing is needed. Moreover, only one service was implemented and tested. Additional services are being considered. Finally, it is necessary for the framework to adhere to HIPAA and GDPR regulations. Some of these regulations, such as access controls and audit trails for the transactions, are already implemented. Encryption of the data both during transfer and at rest is required. Mechanisms for vetting e-health services with respect to regulations compliance are also needed.

Author Contributions: Conceptualization, W.C.H. and Z.B.M.; methodology and formal analysis, W.C.H. and Z.B.M.; software, W.C.H. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Acknowledgments: The authors would like to acknowledge the valuable feedback provided by the reviewers.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Ben-Assuli, O. Electronic health records, adoption, quality of care, legal and privacy issues and their implementation in emergency departments. *Health Policy* **2015**, *119*, 287–297. [CrossRef] [PubMed]
2. The Office of the National Coordinator for Health Information Technology. Health IT Certification Program Overview, 2020. Available online: <https://www.healthit.gov/sites/default/files/PUBLICHealthITCertificationProgramOverview.pdf> (accessed on 26 October 2020).
3. Health and Human Services Department. 2015 Edition Health Information Technology (Health IT) Certification Criteria, 2015 Edition Base Electronic Health Record (EHR) Definition, and ONC Health IT Certification Program Modifications. Federal Register, 2016. Available online: <https://www.federalregister.gov/documents/2015/10/16/2015-25597/2015-edition-health-information-technology-health-it-certification-criteria-2015-edition-base> (accessed on 26 October 2020).
4. Epic Systems Corporation. *Epic*. Available online: <https://www.epic.com/> (accessed on 26 October 2020).
5. Cerner Corporation. *Cerner*. Available online: <https://www.cerner.com/> (accessed on 26 October 2020).
6. Medical Information Technology, Inc. *Meditech*. Available online: <https://ehr.meditech.com/> (accessed on 26 October 2020).
7. HL7 International. *C-CDA (HL7 CDA R2 Implementation Guide: Consolidated CDA Templates for Clinical Notes—US Realm)*. 2019. Available online: http://www.hl7.org/implement/standards/product_brief.cfm?product_id=492/ (accessed on 26 October 2020).
8. Health Level Seven International (HL7). *Index—FHIR v4.0.1*. 2019. Available online: <https://www.hl7.org/fhir/> (accessed on 26 October 2020).
9. Hersh, W.R.; Totten, A.M.; Eden, K.B.; Devine, B.; Gorman, P.; Kassakian, S.Z.; Woods, S.S.; Daeges, M.; Pappas, M.; McDonagh, M.S. Outcomes From Health Information Exchange: Systematic Review and Future Research Needs. *JMIR Med. Inform.* **2015**, *3*, e39. [CrossRef] [PubMed]
10. Williams, C.; Mostashari, F.; Mertz, K.; Hogin, E.; Atwal, P. The Office Of The National Coordinator: The Strategy For Advancing The Exchange Of Health Information. *Health Aff.* **2012**, *31*, 527–536. [CrossRef] [PubMed]

11. Johnson, C.; Pylypchuk, Y.; Patel, V. *Methods Used to Enable Interoperability among U.S. Non-Federal Acute Care Hospitals in 2017*; Office of the National Coordinator for Health Information Technology, 2018. Available online: https://www.healthit.gov/sites/default/files/page/2018-12/Methods-Used-to-Enable-Interoperability-among-U.S.-NonFederal-Acute-Care-Hospitals-in-2017_0.pdf (accessed on 26 October 2020).
12. *eHealth Exchange*. Available online: <https://ehealthexchange.org/> (accessed on 26 October 2020).
13. European Commission. Exchange of Electronic Health Records across the EU. 2020. Available online: <https://ec.europa.eu/digital-single-market/en/exchange-electronic-health-records-across-eu> (accessed on 26 October 2020).
14. Connecting Europe Facility. eHealth Digital Service Infrastructure, 2020. Available online: <https://ec.europa.eu/cefdigital/wiki/display/EHOPERATIONS/eHDSI+Mission> (accessed on 26 October 2020).
15. Office of the National Coordinator for Health Information Technology. *2018 Report to Congress: Annual Update on the Adoption of a Nationwide System for the Electronic Use and Exchange of Health Information*; 2018. Available online: <https://www.healthit.gov/sites/default/files/page/2018-12/2018-HITECH-report-to-congress.pdf> (accessed on 26 October 2020).
16. Eric, H.; Shan, H.; Kevin, I.; Andy, K.; Katherine, L. *A Framework for Cross-Organizational Patient Identity Management*; The Sequoia Project. 2015. Available online: <https://sequoiaproject.org/wp-content/uploads/2015/11/The-Sequoia-Project-Framework-for-Patient-Identity-Management.pdf> (accessed on 26 October 2020).
17. Kruse, C.S.; Argueta, D.A.; Lopez, L.; Nair, A. Patient and Provider Attitudes Toward the Use of Patient Portals for the Management of Chronic Disease: A Systematic Review. *J. Med. Internet Res.* **2015**, *17*, e40. [[CrossRef](#)] [[PubMed](#)]
18. Yocom, C.L. *Health Information Technology: HHS Should Assess the Effectiveness of Its Efforts to Enhance Patient Access to and Use of Electronic Health Information*; U.S. Government Accountability Office. 2017. Available online: <https://www.gao.gov/assets/690/683388.pdf> (accessed on 26 October 2020).
19. Lye, C.T.; Forman, H.P.; Gao, R.; Daniel, J.G.; Hsiao, A.L.; Mann, M.K.; deBronkart, D.; Campos, H.O.; Krumholz, H.M. Assessment of US Hospital Compliance With Regulations for Patients' Requests for Medical Records. *JAMA Netw. Open* **2018**, *1*, e183014, [[CrossRef](#)] [[PubMed](#)]
20. Spil, T.; Klein, R. Personal health records success: Why Google Health failed and what does that mean for Microsoft HealthVault? In Proceedings of the 47th Hawaii International Conference on System Sciences, Waikoloa, HI, USA, 6–9 January 2014; pp. 2818–2827.
21. Tang, P.C.; Ash, J.S.; Bates, D.W.; Overhage, J.M.; Sands, D.Z. Personal health records: definitions, benefits, and strategies for overcoming barriers to adoption. *J. Am. Med. Inform. Assoc.* **2006**, *13*, 121–126, [[CrossRef](#)] [[PubMed](#)]
22. Showell, C. Barriers to the use of personal health records by patients: a structured review. *PeerJ* **2017**, *5*, e3268–e3268, [[CrossRef](#)]
23. Liu, L.S.; Shih, P.C.; Hayes, G.R. Barriers to the Adoption and Use of Personal Health Record Systems. In Proceedings of the iConference, Seattle, WA, USA, 8–11 February 2011; pp. 363–370, [[CrossRef](#)]
24. NORC at the University of Chicago. *Evaluation of the Personal Health Record Pilot for Medicare Fee-For Service Enrollees from South Carolina*; Technical Report; U.S. Department of Health and Human Services: Bethesda, MD, USA, 2010.
25. Schwartz, P.H.; Caine, K.; Alpert, S.A.; Meslin, E.M.; Carroll, A.E.; Tierney, W.M. Patient preferences in controlling access to their electronic health records: A prospective cohort study in primary care. *J. Gen. Intern. Med.* **2015**, *30*, S25–S30, [[CrossRef](#)]
26. Tierney, W.M.; Alpert, S.A.; Byrket, A.; Caine, K.; Leventhal, J.C.; Meslin, E.M.; Schwartz, P.H. Provider responses to patients controlling access to their electronic health records: a prospective cohort study in primary care. *J. Gen. Intern. Med.* **2015**, *30*, S31–S37, [[CrossRef](#)]
27. Caine, K.; Tierney, W.M. Point and Counterpoint: Patient Control of Access to Data in Their Electronic Health Records. *J. Gen. Intern. Med.* **2015**, *30*, 38–41, [[CrossRef](#)]
28. Park, Y.; Yoon, H.J. Understanding Personal Health Record and Facilitating its Market. *Healthc. Inform. Res.* **2020**, *26*, 248–250. [[CrossRef](#)] [[PubMed](#)]
29. Sulieman, L.; Steitz, B.; Rosenbloom, S.T. Analysis of Employee Patient Portal Use and Electronic Health Record Access at an Academic Medical Center. *Appl. Clin. Inform.* **2020**, *11*, 433–441. [[PubMed](#)]

30. Miled, Z.B.; Haas, K.; Black, C.M.; Khandker, R.K.; Chandrasekaran, V.; Lipton, R.; Boustani, M.A. Predicting dementia with routine care EMR data. *Artif. Intell. Med.* **2020**, *102*, 101771. [[CrossRef](#)] [[PubMed](#)]
31. Ali, F.; El-Sappagh, S.; Islam, S.R.; Ali, A.; Attique, M.; Imran, M.; Kwak, K.S. An intelligent healthcare monitoring framework using wearable sensors and social networking data. *Fut. Gen. Comp. Syst.* **2020**, *114*, 23–43. [[CrossRef](#)]
32. Ali, F.; El-Sappagh, S.; Islam, S.R.; Kwak, D.; Ali, A.; Imran, M.; Kwak, K.S. A smart healthcare monitoring system for heart disease prediction based on ensemble deep learning and feature fusion. *Inf. Fus.* **2020**, *63*, 208–222. [[CrossRef](#)]
33. Mettler, T.; Eurich, M. A “design-pattern”-based approach for analyzing e-health business models. *Health Policy Technol.* **2012**, *1*, 77–85. [[CrossRef](#)]
34. Kurose, J.F.; Ross, K.W. *Computer Networking: A Top-Down Approach*, 6 ed.; Pearson: Upper Saddle River, NJ, USA, 2012.
35. Orlando, L.A.; Buchanan, A.H.; Hahn, S.E.; Christianson, C.A.; Powell, K.P.; Skinner, C.S.; Chesnut, B.; Blach, C.; Due, B.; Ginsburg, G.S.; et al. Development and validation of a primary care-based family health history and decision support program (MeTree). *NC Med. J.* **2013**, *74*, 287.
36. Cohn, W.F.; Ropka, M.; Pelletier, S.; Barrett, J.; Kinzie, M.; Harrison, M.; Liu, Z.; Miesfeldt, S.; Tucker, A.; Worrall, B.; et al. Health Heritage[©], a web-based tool for the collection and assessment of family health history: initial user experience and analytic validity. *Public Health Genom.* **2010**, *13*, 477–491. [[CrossRef](#)]
37. Omicini, A.; Denti, E. From tuple spaces to tuple centres. *Sci. Comp. Programm.* **2001**, *41*, 277–294. [[CrossRef](#)]
38. Urovi, V.; Olivieri, A.C.; Bromuri, S.; Fornara, N.; Schumacher, M.I. A Peer to Peer Agent Coordination Framework for IHE Based Cross-community Health Record Exchange. In Proceedings of the 28th Annual ACM Symposium on Applied Computing, Coimbra, Portugal, 18–22 March 2013; pp. 1355–1362. [[CrossRef](#)]
39. King, Z. P2HR, a Personalized Condition-Driven Person Health Record. Doctoral Thesis, Purdue University, West Lafayette, IN, USA, 2017.
40. Kim, J.W.; Lee, A.R.; Kim, M.G.; Kim, I.K.; Lee, E.J. Patient-centric medication history recording system using blockchain. In Proceedings of the IEEE International Conference on Bioinformatics and Biomedicine (BIBM), San Diego, CA, USA, 18–21 November 2019; pp. 1513–1517.
41. Hwang, K.; Dongarra, J.; Fox, G.C. *Distributed and Cloud Computing: From Parallel Processing to the Internet of Things*, 1st ed.; Morgan Kaufmann Publishers Inc.: San Francisco, CA, USA, 2011.
42. Mahy, P. Matthews, J.R. *Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STURN)*; IETF. 2010. Available online: <https://tools.ietf.org/html/rfc5766> (accessed on 26 October 2020).
43. Gray, J. Notes on Data Base Operating Systems. In *Operating Systems, an Advanced Course*; Springer: London, UK, 1978; pp. 393–481.
44. Lamport, L. The Part-time Parliament. *Trans. Comput. Syst.* **1998**, *16*, 133–169, [[CrossRef](#)]
45. Castro, M.; Liskov, B. Practical Byzantine Fault Tolerance and Proactive Recovery. *Trans. Comput. Syst.* **2002**, *20*, 398–461. [[CrossRef](#)]
46. Nakamoto, S. Bitcoin: A peer-to-peer electronic cash system. Technical Report, Manubot, 2019. Available online: <https://git.dhimmel.com/bitcoin-whitepaper/> (accessed on 26 October 2020).
47. Naik, A.R.; Keshavamurthy, B.N. Next level peer-to-peer overlay networks under high churns: A survey. *Peer Peer Netw. Appl.* **2020**, *13*, 905–931. [[CrossRef](#)]
48. Mate, S.; Köpcke, F.; Toddenroth, D.; Martin, M.; Prokosch, H.U.; Bürkle, T.; Ganslandt, T. Ontology-based data integration between clinical and research systems. *PLoS ONE* **2015**, *10*, e0116656. [[CrossRef](#)] [[PubMed](#)]
49. Garde, S.; Knaup, P.; Hovenga, E.J.; Heard, S. Towards semantic interoperability for electronic health records. *Meth. Inf. Med.* **2007**, *46*, 332–343.
50. Spoladore, D. Ontology-based decision support systems for health data management to support collaboration in ambient assisted living and work reintegration. In Proceedings of the Working Conference on Virtual Enterprises, Vicenza, Italy, 18–20 September 2017; pp. 341–352.
51. Spoladore, D.; Sacco, M. Semantic and dweller-based decision support system for the reconfiguration of domestic environments: RecAAL. *Electronics* **2018**, *7*, 179. [[CrossRef](#)]
52. Riaño, D.; Real, F.; López-Vallverdú, J.A.; Campana, F.; Ercolani, S.; Mecocci, P.; Annicchiarico, R.; Caltagirone, C. An ontology-based personalization of health-care knowledge to support clinical decisions for chronically ill patients. *J. Biomed. Inform.* **2012**, *45*, 429–446.

53. Health Level Seven International (HL7). *HL7 FHIR Release 4*. 2019. Available online: <https://www.hl7.org/fhir/rdf.html#ontologies> (accessed on 26 October 2020).
54. World Health Organization. *International Statistical Classification of Diseases and Related Health Problems 10th Revision*. 2016. Available online: <https://icd.who.int/browse10/2016/en> (accessed on 26 October 2020).
55. Haas, K.; Ben Miled, Z.; Mahoui, M. Medication Adherence Prediction Through Online Social Forums: A Case Study of Fibromyalgia. *JMIR Med. Inform.* **2019**, *7*, e12561. [CrossRef]
56. Bender, D.; Sartipi, K. HL7 FHIR: An Agile and RESTful approach to healthcare information exchange. In Proceedings of the 26th IEEE International Symposium on Computer-Based Medical Systems, Porto, Portugal, 20–22 June 2013; pp. 326–331. [CrossRef]
57. Center for Medicare & Medicaid Services. *Blue Button 2.0*, 2019. Available online: <https://bluebutton.cms.gov/> (accessed on 26 October 2020).
58. Mandel, J.C.; Kreda, D.A.; Mandl, K.D.; Kohane, I.S.; Ramoni, R.B. SMART on FHIR: A standards-based, interoperable apps platform for electronic health records. *J. Am. Med. Inform. Assoc.* **2016**, *23*, 899–908, [CrossRef]
59. FIPS PUB 180-4: Secure Hash Standard (SHS). 2015. Available online: <https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.180-4.pdf> (accessed on 26 October 2020).
60. Macedonia, M.R.; Zyda, M.J. A taxonomy for networked virtual environments. *MultiMedia* **1997**, *4*, 48–56.
61. Neumann, C.; Prigent, N.; Varvello, M.; Suh, K. Challenges in Peer-to-peer Gaming. *Comput. Commun. Rev.* **2007**, *37*, 79–82, [CrossRef]
62. Russell, S.; Norvig, P. *Artificial Intelligence: A Modern Approach*, 3 ed.; Pearson: Upper Saddle River, NJ, USA, 2010.
63. Agency for Healthcare Research and Quality. *Medical Expenditure Panel Survey*. 2006. Available online: <https://meps.ahrq.gov/mepsweb/> (accessed on 26 October 2020).
64. Lucas, P. Bayesian Networks in Medicine: A Model-based Approach to Medical Decision Making. 2001. Available online: https://www.researchgate.net/profile/Peter_J_Lucas/publication/2386306_Bayesian_Networks_in_Medicine_a_Model-based_Approach_to_Medical_Decision_Making/links/00b49516dbf4c915ec000000.pdf (accessed on 26 October 2020).
65. Bayes, T. LII. An essay towards solving a problem in the doctrine of chances. By the late Rev. Mr. Bayes, communicated by Mr. Price, in a letter to John Canton, MA and FRS. *Philos. Trans. R. Soc. Lond.* **1763**, *53*, 370–418.
66. Agency for Healthcare Research and Quality. *MEPS HC-181 2015 Medical Conditions*. 2017. Available online: https://meps.ahrq.gov/data_stats/download_data/pufs/h181/h181doc.pdf (accessed on 26 October 2020).
67. Google.com. *Go*, 2019. Available online: <https://golang.org/> (accessed on 26 October 2020).
68. MongoDB Inc. *MongoDB Database*. 2019. Available online: <https://www.mongodb.com> (accessed on 26 October 2020).
69. Niemeyer, G. The MongoDB driver for Go. 2019. Available online: <https://github.com/globalsign/mgo> (accessed on 26 October 2020).
70. Mitre. *Generic FHIR Server Implementation in GoLang*, 2018. Available online: <https://github.com/intervention-engine/fhir> (accessed on 26 October 2020).
71. Fielding, R.T. Architectural Styles and the Design of Network-based Software Architectures. Doctoral Thesis, University of California, Irvine, CA, USA, 2000.

Publisher’s Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).