A Medical Data Reliability Assessment Model

Bandar Alhaqbani¹, Audun Jøsang², and Colin Fidge³

¹ Information Security Institute, Queensland University of Technology, Brisbane, Australia,
b.alhaqbani@isi.qut.edu.au

² UniK Graduate Center, University of Oslo, Oslo, Norway, josang@unik.no

³ Faculty of Science and Technology, Queensland University of Technology, Brisbane, Australia,
c.fidge@qut.edu.au

Received 15 January 2009; received in revised form 25 June 2009; accepted 16 July 2009

Abstract

There is currently a strong focus worldwide on the potential of large-scale Electronic Health Record (EHR) systems to cut costs and improve patient outcomes through increased efficiency. This is accomplished by aggregating medical data from isolated Electronic Medical Record databases maintained by different healthcare providers. Concerns about the privacy and reliability of Electronic Health Records are crucial to healthcare service consumers. Traditional security mechanisms are designed to satisfy confidentiality, integrity, and availability requirements, but they fail to provide a measurement tool for data reliability from a data entry perspective. In this paper, we introduce a Medical Data Reliability Assessment (MDRA) service model to assess the reliability of medical data by evaluating the trustworthiness of its sources, usually the healthcare provider which created the data and the medical practitioner who diagnosed the patient and authorised entry of this data into the patient's medical record. The result is then expressed by manipulating health record metadata to alert medical practitioners relying on the information to possible reliability problems.

Key words: Electronic Health Records, Trustworthiness, Reputation, Reliability, Subjective Logic

1 Introduction

The healthcare domain stands to gain enormously from the increased adoption of Information and Communications Technologies (ICT). Electronic Health Record (EHR) systems are the latest evolution of healthcare ICT, and countries such as Australia, the United Kingdom, and the USA are working on plans for national EHR systems [3].

An Electronic Health Record is defined by Iakovidis [4] as “digitally stored healthcare information about an individual's lifetime with the purpose of supporting continuity of care, education and research, and ensuring confidentiality at all times”. It is a mechanism for integrating healthcare information currently collected in both paper files and Electronic Medical Record (EMR) databases by a variety of healthcare providers [17]. Electronic Health Records (EHRs) enable the efficient communication of medical information, and thus reduce costs and administrative overheads [1], [3]. Furthermore, they aim to provide authorised healthcare professionals with legitimate access to a wealth of historical medical data at one access point.

However, to achieve these potential benefits, the healthcare industry needs to overcome several significant obstacles, in particular having confidence in the reliability of the EHRs' medical data. Concern about data reliability is a crucial factor that may have a strong effect on how medical practitioners use EHRs [4]. This issue arises because the EHR's medical data is composed from different healthcare providers' Electronic Medical Record (EMR) systems and from paper-based medical reports and referrals that patients received from those healthcare providers who do not have an EMR system or an electronic connection to the EHR system. By using the EHR system, a medical practitioner will be exposed to historical medical data with varying levels of reliability; the data might originate from a healthcare provider that does not satisfy patient safety requirements, e.g. is known to habitually enter inaccurate or incomplete data, or be entered by a medical practitioner who fails to satisfy medical safety practices, e.g. is known to violate medical procedures. As a consequence, the reliability of the EHR medical data depends on the trustworthiness of the source and the creator of the captured data. However, in medical practice medical data is usually assumed reliable *a priori* so, in the absence of a reliability evaluation technique, all data will be valued equally. In the following scenario we demonstrate the harmful impact that may occur if a measure of medical data reliability is not incorporated in the EHR system.

Patient 'Sara' from Sydney notices a mole on her back and goes to her General Practitioner 'Tony' who diagnoses the case, informs her that it is not a skin cancer, and performs a small surgery to remove it. Three months later, Sara gets a new job and moves to Cairns. Sara becomes ill and starts losing weight. She visits her new GP 'Frank' who, thanks to the EHR system, looks into her medical history. Seeing nothing relevant in her EHR, he diagnoses the case as depression and prescribes anti-depressants. Six months later, Sara's health gets worse and her GP Frank requests some pathology tests which reveal that Sara has a skin cancer, which has now advanced to a life-threatening state. Frank looks again into Sara's EHR, and specifically the diagnosis made by Tony, and concludes that Tony's diagnosis was not correct. Frank also finds that Tony has had many medical misdiagnosis cases lodged against him in the past two years.

If the EHR system was equipped with a mechanism to evaluate the EHR data's reliability, Frank would have been alerted to Tony's potentially unreliable diagnosis each time he accessed Sara's EHR, allowing him to detect the cancer earlier.

However, current EHR network models (e.g. the Centralised health information model and the Health record data bank model) [19] fail to alert medical practitioners to possible reliability problems with medical data. Therefore, we need a mechanism to alert medical practitioners about possibly unreliable data fields in an Electronic Health Record. This does not imply that such fields should be deleted from the patient's EHR, however. Instead, they should be retained but marked, via appropriate metadata, to indicate their low levels of reliability. Such a reliability evaluation will allow medical practitioners to treat potentially unreliable data with caution, and either test it independently, or check for corroborating evidence in other EHR fields.

In this paper, we propose a Medical Data Reliability Assessment (MDRA) model for validating the reliability of medical data used to construct Electronic Health Records. This is achieved by using metadata to extract information about the medical data's sources (e.g. healthcare providers and medical practitioners). This information is then used in a statistical process to derive a total reliability measure for each EHR data field. The resulting reliability estimate can be expressed in the EHR presented to a medical practitioner to indicate a reliability problem if one exists.

2 Related Work

Reputation systems represent an important input for assessing the trust (or reliability) of a certain agent or service. These systems provide a reputation score for an agent calculated from the agent's ratings as voted on by others who have experienced a transaction with the agent. For instance, eBay's (Site 1) feedback forum is one of the earliest reputation systems; it collects buyers' feedback (either +1, 0, or -1) and aggregates them equally [16] to produce a global reputation score for the seller. The global score is further processed to provide the percentage of positive feedback that is gained by the seller. However, this additive scheme ignores the personalised nature of reputation measures [15]. A slightly better approach, the average reputation scheme [18] provides an improved calculation because it computes the reputation score as the average of all ratings. This principle is used in the reputation systems of many commercial web sites, such as Revolution Health (Site 2) and Amazon (Site 3). Although the average reputation scheme is better than the additive scheme, it still has the same weaknesses.

In the Peer-to-Peer (P2P) research arena, many reputation models have been proposed to assist in assigning reputation scores to those agents within the P2P network. These scores help an agent (service seeker) to make its own decision to trust and connect to the most honest and reliable agents (service providers). EigenTrust [14] is a reputation-based trust management system that aims to minimise malicious behavior in a peer-to-peer network. It computes the agents' trust scores through repeated and iterative multiplication and aggregation of trust scores along transitive chains until the trust scores for all agent members of the P2P community converge to stable values. PeerTrust [22], [23] is another reputation-based trust management system for P2P eCommerce communities. It is even more cautious and examines the received ratings for their quality. It uses five factors to do so, namely feedback in terms of the amount of satisfaction, the number of transactions, the transaction's context factor, and the community context factor. These factors are used to discount the agent's trust value. However, our work differs from these two models in two aspects. Firstly, in the healthcare context, it's crucial to have on hand the identity of the agent who created the medical data (i.e. the healthcare provider or medical practitioner) in order to ensure accountability. In this way, the healthcare context differs significantly from the P2P context. Secondly, in our MDRA model we follow a sounder mathematical basis by using beta and dirichlet probability density functions for combining feedback and for expressing reputation ratings, and subjective logic to represent the trust value where we consider agent's uncertainty factor.

A Bayesian approach is used in more sophisticated reputation systems [20], [21] to produce the reputation score. For example, Wang et al. [15] proposed a reputation model that uses beta probability functions to represent the distribution of trust values according to an interaction history. This model calculates trust either by considering direct observations, if any, or by taking the recommendation of neighbouring agents. However, this model does not distinguish between two trust aspects in its calculation, namely *functional trust* and *referral trust* [10], in which case functional trust values are assigned to neighbors in calculating the transitive trust on a specific agent, while the referral trust supports accurate calculation. TRAVOS [20] is another system that uses the Bayesian approach to calculate a reputation score from binomial ratings and it considers referral trust in its transitive trust calculation. However, in the absence of any evidence, the TRAVOS system will assign an agent a default 0.5 reputation score that results from using the initial settings for the Bayesian parameters α and β . This system does not consider other factors that will have an impact on the default reputation score for these new agents. By contrast, our Medical Data Reliability Assessment model employs a dynamic community base rate that is the average reputation score of the whole community that the agent belongs to. This dynamic community base rate is used in evaluating the reputation of any known or unknown agent, which improves the reliability estimation process since the community base rate dynamically reflects the trustworthiness of the whole community at any one time.

Even more relevant to our research is Hedaquin [2], a system for measuring the quality of health data that is entered into a patient's health record. Hedaquin is based on a Beta reputation system and uses the credentials of the health data supplier, ratings for the health data supplier, and metadata supplied by measuring devices as measures. Hedaquin's goal is similar to our Medical Data Reliability Assessment (MDRA) model, but instead of assessing the quality of the raw health data, our MDRA model assesses the trustworthiness of the medical data as entered into the patient's Electronic Health Record (EHR). Hedaquin uses some ad hoc factors to discount the final quality value and does not provide an accurate trustworthiness estimate of new agents because it follows

the same approach as TRAVOS. By contrast, the approach for assessing data reliability in our MDRA model uses two reputation systems, namely Beta and Dirichlet, which accepts binomial and multinomial ratings and makes the MDRA model capable of expanding and accepting ratings from various trusted agents. In addition, the MDRA model assesses new agents by considering their surrounding community's trustworthiness.

3 Trust and Reputation Systems

Reputation systems collect ratings about users or service providers from members of a community. The reputation system is then able to compute and publish reputation scores about those users and services. Reputation systems use different rating levels, which might be binomial or multinomial. These reputation scores are used to assist in measuring or evaluating the trust or reliability of a certain agent.

In this section, we review the reputation systems that are used in our model, namely Beta and Dirichlet reputation systems, and the Subjective Logic trust model we employ.

3.1 Beta Reputation System

Binomial reputation systems are based on a beta probability function [5], which can be used to represent the probability distribution of binary events, and are therefore called Beta reputation systems. In a reputation calculation process, the Beta reputation system updates its two parameters α and β (Eq. 1) with the recorded observations about a certain agent, to adjust its statistical beta Probability Density Function (PDF). These observations are classified as either positive, r , or negative, s , observations. The α and β parameters can be computed as functions of r and s according to Eq. 1 below,

$$\alpha = r + Wa, \quad \beta = s + W(1 - a), \quad (1)$$

where a expresses the base rate, and W is the weight of the non-informative prior, and normally $W = 2$. The *a posteriori* reputation score is computed as the expected probability, $E(p)$, as defined by:

$$E(p) = \frac{\alpha}{\alpha + \beta}. \quad (2)$$

As an example, let an agent A have 8 positive and 2 negative observations about agent B . Further assume that the base rate a is set to be 0.5. By using Eqs. 1 and 2, the probability expectation value is equal to 0.8. This can be interpreted as saying that the relative frequency of a positive observation in the future is somewhat uncertain, and that the most likely value is 0.8.

3.2 Dirichlet Reputation System

Multinomial Bayesian systems are based on computing reputation scores by statistical updating of Dirichlet Probability Density Functions (PDF), which therefore are called Dirichlet reputation systems [12], [13]. The *a posteriori* (i.e. the updated) reputation score is computed by combining the *a priori* (i.e. previous) reputation score with new ratings.

In Dirichlet reputation systems agents are allowed to rate other agents or services with any value from a set of predefined rating levels, and the reputation scores are not static but will gradually change with time as a function of the received ratings. Initially, each agent's reputation is defined by the base rate reputation. After ratings about a particular agent have been received, that agent's reputation will change accordingly.

Let there be k different discrete rating levels L . This translates into having a state space of cardinality k for the Dirichlet distribution. Let the rating level be indexed by i . The aggregate ratings for a particular agent are stored as a cumulative vector, expressed as:

$$\vec{R} = (\vec{R}(L_i) \mid i = 1 \dots k). \quad (3)$$

This vector can be computed recursively and can take factors such as longevity and the community base rate into account [12]. The most direct way of representing a reputation score for an agent y is to simply aggregate the rating vector $\vec{R}_y$, which represents all relevant previous ratings. The aggregate rating of a particular level i for agent y is denoted by $\vec{R}_y(L_i)$.

For visualisation of reputation scores, the most natural approach is to define the reputation score as a function of the probability expectation values of each rating level. Before any ratings about a particular agent y have been received, its reputation is defined by the common base rate vector $\vec{a}$. As ratings about a particular agent are collected, the aggregate ratings can be computed recursively [12], [13] and the derived reputation scores will change accordingly. Let $\vec{R}$ represent a target agent's aggregate ratings. Then the vector $\vec{S}$, defined by

$$\vec{S}_y : \left(\vec{S}_y(L_i) = \frac{\vec{R}_y(L_i) + W\vec{a}(L_i)}{W + \sum_{j=1}^k \vec{R}_y(L_j)}; i = 1 \dots k \right), \quad (4)$$

is the corresponding multinomial probability reputation score. Parameter W represents the non-informative prior weight, where $W = 2$ is usually the value of choice, but larger values for constant W can be chosen if a reduced influence of new evidence over the base rate is required.

The reputation score $\vec{S}$ can be interpreted like a multinomial probability measure as an indication of how a particular agent is expected to behave in future transactions. It can easily be verified that

$$\sum_{i=1}^k \vec{S}(L_i) = 1. \quad (5)$$

While informative, the multinomial probability representation can require considerable space on a computer screen because multiple values must be visualised. A more compact form can be used to express the reputation score as a single value in some predefined interval. This can be done by assigning a point value v to each rating level L_i , and computing the normalised weighted point estimate score ε .

Assume, for example, that there are k different rating levels with point values $v(L_i)$ evenly distributed in the range $[0, 1]$ according to $v(L_i) = \frac{i-1}{k-1}$. The point estimate reputation score of a reputation $\vec{R}$ is then:

$$\varepsilon = \sum_{i=1}^k v(L_i) \vec{S}(L_i). \quad (6)$$

Such a point estimate in the interval $[0, 1]$ can be scaled to any range, such as 1–5 stars, a percentage or a probability.

Bootstrapping a reputation system to a stable and conservative state is important. In the framework described above, the base rate distribution $\vec{a}$ will define the initial default reputation for all agents. The base rate can, for example, be evenly distributed over all rating levels, or biased towards either negative or positive rating levels. This must be defined when setting up the reputation system in a specific community.

As an example, consider a rating scale with five levels:

- L1: Bad
- L2: Mediocre
- L3: Average
- L4: Good
- L5: Excellent

We assume a default base rate distribution $a = 0.2$. Before any ratings have been received, the multinomial probability reputation score will be represented as in Figure 1(a).

Now assume that 10 ratings are received, where 5 are bad, and 5 are excellent. This translates into the multinomial probability reputation score of Figure 1(b). The point estimate reputation score is calculated by using Eq. 6, and equals 0.5.

Figure 1: Example multinomial probability expectation

3.3 Subjective Logic

Subjective logic [6], [7], [9] is a type of probabilistic logic that explicitly takes uncertainty and belief ownership into account. Arguments in subjective logic are subjective opinions about states in a state space. A binomial opinion applies to a single proposition, and can be represented as a Beta distribution. A multinomial opinion applies to a collection of propositions, and can be represented as a Dirichlet distribution.

Subjective logic defines a trust metric called *opinion* denoted by $\omega_X^A = (\vec{b}, u, \vec{a})$, which expresses the relying party A 's belief over a state space X . Here $\vec{b}$ represents a belief mass vector over the states of X , and u represents an uncertainty mass where $\vec{b}, u \in [0, 1]$ and $\sum \vec{b} + u = 1$. Vector $\vec{a} \in [0, 1]$ represents the base rates over X , and is used for computing the probability expectation value of a state x as

$$E(x) = \vec{b}(x) + \vec{a}(x)u, \quad (7)$$

meaning that $\vec{a}$ determines how uncertainty contributes to $E(x)$. Binomial opinions are expressed as $\omega_x^A = (b, d, u, a)$ where d denotes disbelief in statement x . For instance, given the statement x : "*David is honest and reliable*", then the opinion can be interpreted as evaluation trust in David. More specifically, the trust target is David, and the trust scope is σ : "*To be honest and reliable*", so that $x \equiv D(\sigma)$. The opinion can be denoted with explicit attributes as $\omega_{D(\sigma)}^A$, but the trust scope can be omitted when it is obvious.

As an example, assume that Alice needs treatment for her knee and asks her GP Bob to recommend a good physiotherapist. When Bob recommends David, Alice would like to get a second opinion, so she asks Claire for her opinion about David. The trust scope in this case can be expressed as σ : “To be a competent physiotherapist”. This situation is illustrated in Figure 2 where the indexes on arrows indicate the order in which the opinions are formed.

Figure 2: Deriving trust from parallel transitive chains

When trust and referrals are expressed as subjective opinions, each transitive trust path Alice→Bob→David, and Alice→Claire→David can be computed with the *transitivity operator*, also called the discounting operator, where the idea is that the referrals from Bob and Claire are discounted as a function of Alice's trust in Bob and Claire respectively. Finally the two paths can be combined using the cumulative or averaging fusion operator. These operators form part of *Subjective Logic* [7], [9], and semantic constraints must be satisfied in order for the transitive trust derivation to be meaningful [8]. Opinions can be uniquely mapped to beta PDFs, and in this sense the fusion operator is equivalent to Bayesian updating. This model is thus both belief-based and Bayesian.

Algebraically, a trust relationship between A and B is denoted $[A, B]$, transitivity of two arcs is indicated using a binary “:” operator, and the fusion of two parallel paths is indicated with a “ $\diamond$ ” operator. The trust network of Figure 2 can then be expressed as:

$$[A, D] = ([A, B] : [B, D]) \diamond ([A, C] : [C, D]). \quad (8)$$

The corresponding transitivity operator for opinions is denoted as “ $\otimes$ ” and the corresponding fusion operator as “ $\oplus$ ”. The mathematical expression for combining the opinions about the trust relationships of Figure 2 is then:

$$\omega_D^A = (\omega_B^A \otimes \omega_D^B) \oplus (\omega_C^A \otimes \omega_D^C). \quad (9)$$

4 Medical Data Reliability Network Structure

Figure 3 shows our proposed network structure for deriving a healthcare provider's level of trust in Electronic Health Record medical data fields, via our Medical Data Reliability Assessment model. In this section, we explain the functionality of each component. The protocol by which these components interact is described in Section 5.

4.1 Healthcare Authority (HA)

A Healthcare Authority is a legal body that records information (metadata) gathered from public sources including, but not limited to, reports received from healthcare providers and medical practitioners about incorrect medical data or procedures, and medical misconduct, non-safety, or malpractice cases. The subject of this information is either a healthcare provider, a medical practitioner, or both. The HA uses this information to produce a ratings vector, in which it ranks each reported case according to its severity. This process can be done by applying previously-defined classification rules to each case.

Figure 3: Medical Data Reliability Network Structure

In addition the HA assigns a base rate for each severity rating level and for the prior behavior of the healthcare agent (either a healthcare provider or medical practitioner). The HA's ratings vector will have k levels representing the severity (danger) levels for reported cases. Here we assume that level 1 denotes the highest level of severity and level $k - 1$ the lowest. Level k denotes the special case of assumed 'perfect' behaviour, in the sense that there have been no adverse reports.

From this, the HA provides authorised or registered healthcare providers with its opinion (arrow 2 in Figure 3) about the medical conduct and practice of a certain healthcare provider or medical practitioner. Here, the HA acts as a Dirichlet reputation system and expresses its trust using the Subjective Logic trust metric *opinion*. For example, $\omega_A^{HA} = (\vec{b}_A^{HA}, u_A^{HA}, \vec{d}_C^{HA})$ is Healthcare Authority HA 's trust opinion about healthcare agent A . $\vec{d}_C^{HA}$ represents HA 's base rate for A 's community (C).

In order to represent the Healthcare Authority's opinion as a single value, it uses a point estimate representation. Due to the fact that the rating levels are not evenly distributed, the HA should manually define point value vector $\vec{m}$ to express its weight for each rating level. Therefore, the reputation score is computed as:

$$\varepsilon_X^{HA} = \sum_{i=1}^k \vec{m}(L_i) \vec{S}(L_i). \quad (10)$$

4.2 Reputation Centre (RC)

A Reputation Centre receives ratings from community members (e.g. patients) about healthcare providers and medical practitioners. These ratings are used by the reputation centre to derive a reputation score for those rated agents; this reputation score represents the RC's subjective opinion (arrow 4 in Figure 3). These opinions can be communicated to healthcare providers. The RC acts as a Dirichlet reputation centre and expresses its opinions in the same way that the Healthcare Authority HA does. For example, $\omega_A^{RC} = (\vec{b}_A^{RC}, u_A^{RC}, \vec{d}_C^{RC})$ is Reputation Centre RC 's opinion about healthcare agent A .

4.3 Medical Data Reliability Assessment (MDRA) Service

The Medical Data Reliability Assessment service is employed by the Electronic Health Record (EHR) system used by a Healthcare Provider (HP) to measure the reliability of medical data sourced from other healthcare agents (e.g. healthcare providers or medical practitioners). The HP has a database that records the HP's interaction experiences with other healthcare agents. These experiences are created from the internal reports that are received from the HP's medical practitioners about those received external medical data. The HP uses this information to either record good or bad experiences with the agents who created these data. The HP's experiences are then used by the MDRA that acts as a Beta reputation system to compute the HP's opinion about a certain healthcare agent. Healthcare provider HP 's opinion about a healthcare agent A (arrow 1 in Figure 3) is denoted as $\omega_A^{HP*} = (b_A^{HP*}, d_A^{HP*}, u_A^{HP*}, a_C^{HP*})$.

In addition, the MDRA can communicate with the HA and the RC to get their opinion about a certain agent, to use in the MDRA's reliability calculation process. Also, the MDRA maintains dynamic opinions about the HA and the RC (arrows 3 and 5 in Figure 3) that are calculated based on opinion comparison [11].

5 MDRA Protocol

Section 4 introduced the components needed to implement our reliability model. In this section we define the protocol whereby these components interact with one another.

The Medical Data Reliability Assessment (MDRA) service is a supporting service for an Electronic Health Record (EHR) system. It is responsible for assessing the reliability of given medical data and then communicating this information to the EHR system to update the medical metadata displayed. This process starts by receiving medical data from the EHR system, then the MDRA starts its investigation by consulting the EHR reputation system and seeks, if necessary, opinions from known parties. In this setup, we assume that each entity, including healthcare providers, medical practitioners, the Healthcare Authority, and the Reputation Centre has a well-defined identity that can be verified in a secure context.

To better understand the functionality of the MDRA service, we use the following steps to depict the messages received and sent by the MDRA service in order to accomplish its task (Figure 4).

1. Healthcare provider HP 's EHR system receives medical data for patient p from healthcare provider $HP2$.
2. The EHR system sends this medical data to the MDRA service to evaluate its reliability.
3. The MDRA extracts medical metadata to identify the source healthcare provider src , who is $HP2$ in this case, and the identity of the medical practitioner mp who created this record.
4. The MDRA accesses the EHR reputation database to find historical interaction experiences between HP and src , and HP and mp .
5. If the recorded experiences do not satisfy HP 's confidence criteria (see Section 6) then:
 - (a) The MDRA requests opinions about src and mp from Healthcare Authority HA .
 - (b) The MDRA requests opinions about src and mp from Reputation Centre RC .
6. The MDRA service uses the received information to calculate the medical data reliability score.
7. The MDRA service sends the result back to the EHR system.
8. The EHR system updates the medical data displayed to reflect the computed reliability score, e.g. by highlighting potentially unreliable data.

Figure 4: Medical data reliability evaluation message sequencing

6 Measuring the Reliability of Medical Data

Assume a healthcare provider *HP* has received medical data about a specific patient. This medical data consists of medical data fields, and each field *MF* has attached metadata. This metadata provides information about the identity of the healthcare provider *src* who produced the data, and of the medical practitioner *mp* who diagnosed the patient and authorised entry of this data into the patient's medical record. In order to evaluate the reliability ϕ_{MF}^{HP} of a given medical data field, *HP*'s MDRA service conducts a trust assessment for those agents responsible for producing each data field *MF*. This process starts by evaluating healthcare provider *HP*'s opinion ω_{src}^{HP} of the source healthcare provider *src*, and *HP*'s opinion ω_{mp}^{HP} of medical practitioner *mp*. Afterwards, the MDRA uses this information to compute the reliability of the medical data by using the fusion operator as in the following equation.

$$\phi_{MF}^{HP} = E(\omega_{src}^{HP} \oplus \omega_{mp}^{HP}) \quad (11)$$

The Electronic Health Record (EHR) system uses the resulting reliability score ϕ_{MF}^{HP} to update medical data field *MF*'s metadata to reflect this score in order to alert the medical practitioner relying on the data if its reliability is low.

In order to compute the reliability of medical data field *MF*, the MDRA needs to evaluate the trustworthiness of its sources *src* and *mp*. However, the MDRA's approach for calculating these values is similar; therefore we will denote the trust target as agent *X* which represents either *src* or *mp*. The MDRA follows two approaches in calculating the trust of a given agent *X* and this approach is determined by evaluating certain criteria which we call the *confidence criteria*. In our system, we define the confidence criteria as the number of interaction experiences *n* had with agent *X* in a period of time *t*. Based on these criteria, the MDRA will use its internal assessment service (Section 6.1) if *HP*'s interaction experiences with agent *X* within time period *t* are greater than or equal to *n*, otherwise it will use an external assessment service (Section 6.2) from which it will seek the healthcare authority and reputation centre's opinions about *X*.

6.1 Internal Assessment

The Medical Data Reliability Assessment service uses healthcare provider HP 's reputation database to derive HP 's opinion ω_X^{HP*} about agent X . Provider HP 's opinion parameters are calculated as follows:

$$\begin{aligned} b_X^{HP} &= \frac{r_X^{HP}}{(r_X^{HP} + s_X^{HP} + 2)} \\ d_X^{HP} &= \frac{s_X^{HP}}{(r_X^{HP} + s_X^{HP} + 2)} \\ u_X^{HP} &= \frac{2}{(r_X^{HP} + s_X^{HP} + 2)} \\ a_C^{HP} &= HP\text{'s base rate for agent } X\text{'s community } C \end{aligned} \quad (12)$$

The base rate in Eq. 12 helps the healthcare authority to set *a priori* trust about a certain agent in the absence of any interaction experiences. On start-up of the reputation system this value is usually set by the authority who provides the reputation system.

Most previous work treats the base rate as a static value that does not change over time. However, this is inadequate for our purposes because the base rate should reflect the evaluator's belief at a certain time towards its targeted community. Therefore, we use healthcare provider HP 's base rate for agent X 's community (the relevant healthcare provider's or medical practitioner's community base rate) to represent the base rate in HP 's opinion. In order to calculate the community base rate, the MDRA aggregates the community's reputation values at time t .

$$\begin{aligned} r_{C,t}^{HP} &= \sum_{M \in C} r_M^{HP} \\ s_{C,t}^{HP} &= \sum_{M \in C} s_M^{HP} \end{aligned} \quad (13)$$

Afterwards, we define the community base rate at time $t + 1$ as the community's reputation score at time t .

$$a_{X,(t+1)}^{HP} = \frac{r_{C,t}^{HP} + r_0}{r_{C,t}^{HP} + s_{C,t}^{HP} + r_0 + s_0} \quad \text{where} \quad \begin{cases} X \in C \\ r_0 = W a_{C,t}^{HP} \\ s_0 = W (1 - a_{C,t}^{HP}) \end{cases} \quad (14)$$

Here r_0 and s_0 represent the *a priori* reputation of community C . They are expressed as a non-informative constant weight W that is distributed over all possible outcomes as a function of the base rate.

Once the MDRA has computed HP 's opinion about X and HP 's experiences with X satisfy the confidence criteria, the MDRA publishes HP 's opinion about X as follows.

$$\omega_X^{HP} = \omega_X^{HP*} \quad (15)$$

Once the MDRA finishes computing HA 's opinions about src and mp , it substitutes these values into Eq. 11 to derive HP 's subjective reliability measure ϕ_{MF}^{HP} on the medical data.

6.2 External Assessment

In this approach, healthcare provider HP seeks opinions from external parties to be combined with HP 's self-computed opinion ω_X^{HP*} in order to derive HP 's overall opinion ω_X^{HP} about agent X . There are two sources of information: relevant Healthcare Authority HA and Reputation Centre RC . Each source sends its opinion about X to HP through a secure communications channel. However, these opinions are discounted by HP 's opinion

about each source. The following equation uses the Subjective Logic fusion operator to compute HP 's opinion about X using HP 's self-computed opinion, HA 's opinion, and RC 's opinion about X .

$$\omega_X^{HP} = \omega_X^{HP*} \oplus \omega_X^{HP:RC} \oplus \omega_X^{HP:HA} \quad (16)$$

The computation process for discounted opinions $\omega_X^{HP:RC}$ and $\omega_X^{HP:HA}$ is the same. Therefore, in the following, we show how to compute $\omega_X^{HP:RC}$, and the same process can be applied to produce $\omega_X^{HP:HA}$.

Firstly, the MDRA needs to compute HP 's opinion about RC , so the MDRA uses an opinion comparison approach via the operator " $\downarrow$ " [11], which here is HP 's distrust of RC , proportional to the greatest difference in point estimates between HP and RC opinions about certain agent. The MDRA selects an agent Z from HP 's database, calculates HP 's opinion ω_Z^{HP*} about Z , and compares it to RC 's opinion ω_Z^{RC} about Z . The following equation shows how HP 's opinion about RC is computed.

$$\omega_{RC}^{HP} = \omega_Z^{HP} \downarrow \omega_Z^{RC} \quad \text{where} \quad \begin{cases} d_{RC}^{HP} = \left| \left(\frac{r_Z^{HP} + 1}{r_Z^{HP} + s_Z^{HP} + 2} \right) - \varepsilon(\bar{R}_Z^{RC}) \right| \\ u_{RC}^{HP} = \max [u_Z^{HP}, u_Z^{RC}] \\ b_{RC}^{HP} = 1 - b_{RC}^{HP} - u_{RC}^{HP} \end{cases} \quad (17)$$

Secondly, the Reputation Centre RC needs to convert its multinomial aggregate ratings $\bar{R}_X^{RC}$ into a binomial opinion. Reputation Centre RC uses the following equation to derive the binomial rating parameters r and s .

$$\begin{aligned} r &= \varepsilon_X^{RC} \sum_{i=1}^k \bar{R}_X^{RC}(x_i) \\ s &= \sum_{i=1}^k \bar{R}_X^{RC}(x_i) - r \end{aligned} \quad (18)$$

Afterwards, RC uses Eq. 12 to derive its binomial parameters b , d , and u . The base rate a parameter is computed as in Eq. 19.

$$a_X^{RC} = \frac{\varepsilon_X^{RC} - b_X^{RC}}{u_X^{RC}} \quad (19)$$

Finally, the MDRA uses ω_{RC}^{HP} and ω_X^{RC} with the transitivity operator to compute the discounted opinion $\omega_X^{HP:RC}$ as shown in the following equation.

$$\omega_X^{HP:RC} = \omega_{RC}^{HP} \otimes \omega_X^{RC} \quad \text{where} \quad \begin{cases} b_X^{HP:RC} = (b_{RC}^{HP} + a_{RC}^{HP} u_{RC}^{HP}) b_X^{RC} \\ d_X^{HP:RC} = (b_{RC}^{HP} + a_{RC}^{HP} u_{RC}^{HP}) d_X^{RC} \\ u_X^{HP:RC} = 1 - b_X^{HP:RC} - d_X^{HP:RC} \\ a_X^{HP:RC} = a_X^{RC} \end{cases} \quad (20)$$

Once the MDRA has computed discounted opinions $\omega_X^{HP:HA}$ and $\omega_X^{HP:RC}$, it uses these values in Eq. 16 to derive HP 's opinion about X , which is either healthcare provider src or medical practitioner mp . Finally, the MDRA substitutes HP 's computed opinion ω_{src}^{HP} of src and ω_{mp}^{HP} of mp into Eq. 11 to derive HP 's subjective reliability measure ϕ_{MF}^{HP} of medical data field MF .

7 Motivational Scenario Revisited

In this section we revisit the motivational scenario from Section 1 to see how our Medical Data Reliability Assessment (MDRA) service would alert medical practitioner Frank to Tony's unreliable medical data. Assume that each system, including healthcare provider Frank's medical system HP , the nationwide healthcare authority HA , and

(a) HA's Reputation System				
Agent	$\vec{R}$	$\vec{a}$		
Tony (<i>T</i>)	(3, 5, 4, 2, 0)	(0.002, 0.006, 0.008, 0.01, 0.974)		
Medical Centre (<i>MC</i>)	(1, 2, 2, 5, 0)	(0.002, 0.004, 0.008, 0.01, 0.976)		
Karen (<i>K</i>)	(0, 0, 0, 1, 0)	(0.002, 0.006, 0.008, 0.01, 0.974)		

(b) HP's Reputation System			(c) RC's Reputation System		
Agent	<i>r</i>	<i>s</i>	Agent	$\vec{R}$	$\vec{a}$
Tony (<i>T</i>)	0	1	Tony (<i>T</i>)	(7, 4, 2, 3, 1)	(0.1, 0.2, 0.3, 0.2, 0.2)
Medical Centre (<i>MC</i>)	2	3	Medical Centre (<i>MC</i>)	(3, 2, 2, 6, 2)	(0.1, 0.2, 0.2, 0.3, 0.2)
Karen (<i>K</i>)	6	1	Karen (<i>K</i>)	(0, 0, 3, 4, 3)	(0.1, 0.2, 0.3, 0.2, 0.2)

Table 1: Reputation scores for the motivational scenario

the government-run reputation centre *RC*, has reputation scores about Tony and his healthcare centre as shown in Table 1. Authority *HA*'s reputation system maintains rating vector $\vec{R}$ (Table 1(a)) that has five elements, four represent severity rating levels (extreme, high, medium, low) and the fifth element represents the well behaviour, the base rate vector $\vec{a}$ represents the priori base rate for those elements in $\vec{R}$. Healthcare Provider *HP*'s reputation system (Table 1(b)) has two values, positive observations *r* and negative observations *s*. Reputation centre *RC* maintains ratings vector $\vec{R}$ (Table 1(c)) which has five rating levels (bad, mediocre, average, good, excellent), the base rate $\vec{a}$ represents the *a priori* base rate for elements in $\vec{R}$.

When Frank requests Sara's EHR, his medical system will evaluate the reliability of each data field. Here we show how the MDRA will evaluate the reliability of Tony's medical data. The MDRA starts by checking its reputation system, where there are insufficient experience entries recorded either with Tony or his medical centre. Therefore, the MDRA requests *HA*'s and *RC*'s opinion about Tony and his medical centre. Healthcare authority *HA* uses its point values $\vec{m} = (0, 0.2, 0.43, 0.67, 1)$ with $\vec{S}$, which is computed using Eq. 4, in Eq. 10 to produce its reputation scores ε_T^{HA} and ε_{MC}^{HA} . Then *HA* uses Eqs. 18, 12 and 19 to compute its binomial opinions $\omega_T^{HA} = (0.33, 0.55, 0.13, 0.38)$ and $\omega_{MC}^{HA} = (0.49, 0.33, 0.18, 0.6)$. Afterwards, the MDRA uses Eq. 17 to compute its opinion about *HA*, with Karen as the subject of this process. The resulting opinion $\omega_{HA}^{HP} = (0.23, 0.1, 0.67, 0.8)$, where the base rate trust has been set to $a_{HA}^{HP} = 0.8$, is used with each of ω_T^{HA} and ω_{MC}^{HA} in Eq. 20 to compute the MDRA's discounted opinion that *HA* holds about Tony, which is $\omega_T^{HP:HA} = (0.25, 0.42, 0.33, 0.33)$, and about Tony's healthcare centre which is $\omega_{MC}^{HP:HA} = (0.37, 0.25, 0.38, 0.6)$.

The MDRA follows the previous approach to compute its discounted opinion held by *RC* about Tony and his medical centre. However, the only difference in this process is the way that *RC* computes its reputation score. Since *RC*'s rating levels are evenly distributed, it uses Eq. 6 to compute its reputation score ε . As a result, the MDRA's discounted opinion for *RC* about Tony is $\omega_T^{HP:RC} = (0.28, 0.55, 0.17, 0.33)$ and about Tony's medical centre is $\omega_{MC}^{HP:RC} = (0.52, 0.4, 0.09, 0.65)$.

In the next step, the MDRA substitutes its internal opinions: $\omega_T^{HP*} = (0, 0.33, 0.67, 0.29, 0.5)$ and $\omega_{MC}^{HP*} = (0.29, 0.43, 0.29, 0.5)$, and its calculated discounted opinions into Eq. 16 to compute *HP*'s opinion about Tony and his medical centre which is $\omega_T^{HP} = (0.24, 0.55, 0.12, 0.17)$ and $\omega_{MC}^{HP} = (0.48, 0.39, 0.06, 0.62)$.

Finally, the MDRA computes *HP*'s reliability measure ϕ about Tony's medical entry by substituting ω_T^{HP} and ω_{MC}^{HP} into Eq. 11 which results in $\phi = 0.018$ which implies that the medical data is unreliable.

The MDRA communicates this result to the EHR system, which then updates the medical metadata for the fields derived from Tony's diagnosis to indicate their low reliability. Now GP Frank will notice the highlighted medical data, which implies that this information is unreliable, and will reinvestigate the case again which will result in discovering the skin cancer in an earlier stage.

8 Conclusion and Future Work

An Electronic Health Record (EHR) system overcomes the problems and limitations that are associated with paper based systems and isolated Electronic Medical Record (EMR) systems; however, it is hindered by several 'soft' security problems such as concerns over privacy and reliability (trust). Medical data reliability is a crucial requirement that can affect a medical practitioner's decisions when using historical medical data. In the current situation, all historical medical data are trusted equally; however, this should not be the case.

In this paper, we presented a Medical Data Reliability Assessment (MDRA) model that can augment an EHR with metadata containing a reliability measure of medical data. In this process, several reputation systems have been used together with a subjective logic trust analysis to produce a reliability measure. The MDRA can communicate this reliability score to the EHR displayed on the medical practitioner's computer to alert the medical practitioner to any reliability concerns.

In future work we will demonstrate this model by implementing it in a workflow system, based on the YAWL (Site 4) workflow system, so that the displayed data responds to each reliability assessment request that comes from any healthcare business process. Also, we will use the extended attributes implemented in YAWL to reflect the calculated reliability score of given medical data.

In this paper, we assumed that the healthcare authority (HA) maintains a ratings vector about healthcare agents, but we did not discuss how the HA treats previously investigated and dismissed cases of medical misconduct in its ratings calculation. Therefore, in future work we will develop an appropriate healthcare-oriented time-variant processing mechanism which will allow the HA to produce ratings scores which exonerated medical practitioners to recover their reputation.

Acknowledgements

We wish to thank the anonymous reviewers for their many helpful comments and suggestions. The first author's work is sponsored by King Saud bin Abdulaziz University for Health Sciences – Saudi Arabia.

Websites List

Site 1: Ebay

<http://www.ebay.com>

Site 2: Revolution Health

<http://www.revolutionhealth.com>

Site 3: Amazon

<http://www.amazon.com>

Site 4: YAWL

<http://www.yawl-system.com>

References

- [1] B. Blobel. Authorisation and Access Control for Electronic Health Record Systems. International Journal of Medical Informatics Realizing Security into the Electronic Health Record, vol. 73, no. 3, pp. 251–257, 2004.
- [2] T. Deursen, P. Koster, and M. Petkovic. Hedaquin: A Reputation-based Health Data Quality Indicator. In Proceedings of the 3rd International Workshop on Security and Trust Management (STM 2007), February 2008, volume 197 of Electronic Notes in Theoretical Computer Science, pp. 159–167, 2008.

- [3] T. Gunter and N. Terry. The Emergence of National Electronic Health Record Architectures in the United States and Australia: Models, Costs, and Questions. *Journal of Medical Internet Research*, vol. 7, no. 1, 2005.
- [4] I. Iakovidis. Towards Personal Health Record: Current Situation, Obstacles and Trends in Implementation of Electronic Healthcare Record in Europe. *International Journal of Medical Informatics*, vol. 52, no. 1-3, pp. 105-115, 1998.
- [5] R. Ismail and A. Jøsang. The Beta Reputation. In *Proceedings of the 15th Bled Conference on Electronic Conference*, 2002.
- [6] A. Jøsang. Artificial Reasoning with Subjective Logic. In Abhaya Nayak and Maurice Pagnucco, editors, *Proceedings of the 2nd Australian Workshop on Commonsense Reasoning*, Perth, December 1997, volume 65 of CRPIT. Australian Computer Society, 1997.
- [7] A. Jøsang. A Logic for Uncertain Probabilities. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 9, no. 3, pp. 279-212, 2001.
- [8] A. Jøsang. Semantic Constraints for Trust Transitivity. In Sven Hartmann and Markus Stumptner, editors, *Proceedings of the Asia-Pacific Conference of Conceptual Modeling (APCCM)*, Newcastle, Australia, February 2005, volume 43 of CRPIT. Australian Computer Society, 2005.
- [9] A. Jøsang. Probabilistic Logic under Uncertainty. In Joachim Gudmundsson and C. Barry Jay, editors, *Proceedings of Computing: The Australian Theory Symposium (CATS2007)*, Ballarat, Australia, January 2007, volume 65 of CRPIT. Australian Computer Society, 2007.
- [10] A. Jøsang. Trust and Reputation Systems. In A. Aldini and R. Gorrieri, editors, *Foundations of Security Analysis and Design IV*, volume LNCS 4677. Springer, 2007.
- [11] A. Jøsang, T. Bhuiyan, and C. Cox. Combining Trust and Reputation Management for Web-Based Services. In *Proceedings of the 5th International Conference on Trust, Privacy and Security in Digital Business (TrustBus2008)*, Turin, September, 2008.
- [12] A. Jøsang and J. Haller. Dirichlet Reputation Systems. In *Proceedings of the 2nd International Conference on Availability, Reliability and Security (ARES 2007)*, Vienna, Austria, April, 2007.
- [13] A. Jøsang, X. Luo, and X. Chen. Continuous Ratings in Discrete Bayesian Reputation Systems. In *Proceedings of the Joint iTrust and PST Conferences on Privacy, Trust Management and Security (IFIPTM 2008)*, Trondheim, June, 2008.
- [14] S. Kamvar, M. Schlosser, and H. Garcia-Molina. The Eigentrust Algorithm for Reputation Management in P2P Networks. In *WWW '03: Proceedings of the 12th international conference on World Wide Web*, pp. 640-651, New York, NY, USA, 2003. ACM.
- [15] L. Mui, M. Mohtashemi, C. Ang, P. Szolovits, and A. Halberstadt. Rating in Distributed Systems: A Bayesian Approach. In *Proceedings of the Workshop on Information Technologies and Systems (WITS)*, 2001.
- [16] P. Resnick and R. Zeckhauser. Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System. In Michael R. Baye, editor, *The Economics of the Internet and E-commerce*, volume 11 of *Advances in Applied Microeconomics*, pp. 127 - 157. Elsevier Science, 2002.
- [17] W. Rishel, T. Handler, and J. Edwards. A Clear Definition of the Electronic Health Record. Technical report, Gartner, 2005.
- [18] J. Schneider, G. Kortuem, J. Jager, S. Fickas, and Z. Segall. Disseminating Trust Information in Wearable Communities. vol. 4, pp. 245-248, London, UK, 2000. Springer-Verlag.
- [19] A. Shabo, P. Vortman, and B. Robson. Who's Afraid of Lifetime Electronic Medical Records? In *Towards Electronic Health Records (TEHRE 2001)*, 2001.
- [20] W. Teacy, J. Patel, N. Jennings, and M. Luck. TRAVOS: Trust and Reputation in the Context of Inaccurate Information Sources. *Autonomous Agents and Multi-Agent Systems*, vol. 12, no. 2, pp. 183-198, 2006.
- [21] Y. Wang, V. Cahill, E. Gray, C. Harris, and L. Liao. Bayesian Network Based Trust Management. In Laurence Tianruo Yang, Hai Jin, Jianhua Ma, and Theo Ungerer, editors, *Proceedings the Third International Conference in Autonomic and Trusted Computing (ATC)*, Wuhan, China, September 2006, volume 4158 of *Lecture Notes in Computer Science*, pp. 246-257. Springer, 2006.
- [22] L. Xiong and L. Liu. A Reputation-Based Trust Model for Peer-to-Peer eCommerce Communication. In the *IEEE international Conference on E-Commerce (CEC'03)*, 2003.
- [23] L. Xiong and L. Liu. PeerTrust: Supporting Reputation-based Trust for Peer-to-Peer Electronic Communities. *Knowledge and Data Engineering, IEEE Transactions on*, vol. 16, no. 7, pp. 843-857, July 2004.