Informatics for Health: Connected Citizen-Led Wellness and Population Health
R. Randell et al. (Eds.)

doi:10.3233/978-1-61499-753-5-476

Computable Information Governance Contracts

James CUNNINGHAM [a,1], Gary LEEMING [b] and John AINSWORTH [a]
[a] *The Health eResearch Centre and The Farr Institute, Division of Informatics, Imaging and Data Sciences, School of Health Sciences, The University of Manchester*
[b] *Greater Manchester Academic Health Science Network*

Abstract. The risks of relinquishing control of electronic healthcare data for re-use in research are mitigated by the use of data sharing agreements and information governance procedures. These exist as legal, or quasi-legal, textual documents exchanged between data owners. Their existence outside of the digital realm leads to a situation where breaches of an agreement can only be detected and acted on post-hoc. We introduce the design of a system of computable contracts, specified formally, that can enforce the rules of data sharing agreements within the bounds of electronic health care systems.

Keywords. Information Storage and Retrieval, Database Management Systems

1. Introduction

Historically medical records have been paper-based, with the UK only recently moving towards an electronic representation [1], mirroring efforts and trends in Europe, North America and Australia amongst others [2]. Whilst the increased benefits this gives to health care provision and large scale medical research are clear [3] a sense of ownership has grown-up around the physical, paper-based, records an their corresponding electronic form. Value exists to both the organization and the practitioner in maintaining sole ownership of patient data. Economic value is created through the retention of a patient and even within the UK's relatively unified National Health Service (NHS) there exists competition between healthcare providers [4]. The challenges surrounding the reuse and trustworthiness of healthcare data are well recognised [5]. Whilst the use of electronic healthcare data is crucial for research, its continued use for such purposes is grounded in the maintenance of strict ethical frameworks of control [6]. Situations in which these ethical frameworks are perceived as ineffectual can lead to a breakdown in public trust in EHR systems, and in turn will threaten to bring down tighter regulatory controls on the use of data, threatening its effective use moving forwards [7]. Information Governance (IG), the policies procedures and controls that ensure the correct use of data, is therefore a crucial aspect in the use and practice of the digitisation of healthcare data. Without correct IG procedures, that both bring confidence in data use and ensure that this use falls within existing regulatory frameworks, the future benefits of digitised healthcare data will be

[1] Corresponding Author, Email: james.a.cunningham@manchester.ac.uk

severely limited. The documented procedures and legal agreements that specify the correct ethics-legal use of electronic healthcare data are specified and exchanged as documents distinct from the digitised data that they pertain to. There has been no 'big bang' move from the paper to the digital realm whereby all systems were switched at once to a modern representation; the transformation has been slow and piecemeal – an inevitable consequence of the crucial role that patient records play in day-to-day care and the need to maintain the smooth running of healthcare system. As such there are aspects of healthcare that have not yet started the move to the digital, information governance agreements and their use being one such aspect of the domain.

In this paper we describe a framework for specifying information governance contracts in computable form. These computable information governance contracts allow data owners to precisely specify the conditions under which certain actions (such as querying and retrieval) can be performed on medical data. These contracts can then produce electronically signed warrants which grant permissions to users and organisations the right to use data for purposes controlled by the underlying contracts. By attaching these warrants to formal requests for data both the data controller and the data user are provided with evidence that can ensure the correct and verifiable delivery of data for the purpose that its provision was intended, although production and delivery of the data itself lies outside the scope of this framework.

2. Method

The system we describe is based at its core on the notion of **contracts** for specifying the scope of allowed behaviour in terms of data sharing within an EHR system and **warrants** that specify data that can be returned from a system for a particular data access request. Contracts are mapped from agreements and exchanged between data providers. For a given access request for data a **request** object is specified which describes who is requesting the data and for what purpose that data will be used. Given a request and a contract a warrant is produced that describes the data accessible by that user for that purpose from a given data source. This warrant is then provided by the user to the data owner and used to determine what data is produced for that user.

The system as described below is given as a series of functional data types specified using the semantics of the F# programming language. Briefly, a type has an identifying name and a series of constructors (separated by the | symbol) that are used to define a value of that type. A constructor can, if needed, be given a series of values or the types specified by the 'of' keyword and separated by `*`. The types used to construct the warrants and contracts used by the system are given in the following.

A `Purpose` describes the use to which data will be put. Initially we have identified the core categories of `PatientCare` – accessing data in order to directly influence the course of care of a patient, `Research` – secondary research use, and `Benchmarking` – general data auditing. Additionally a `Purpose` can be defined as being `Any`, which indicates that once data is received it will potentially be used for any purpose within any standard constraints.

```
type Purpose
| Any
```

```
| PatientCare
| Research
| Benchmarking
```

An `Action` describes something taking place. In the simple case an `Action` can be For a `Purpose` as described above. An `Action` can also be specified as taking place `Until` or `After` a given date and time. Further a `Choice` represents a means of grouping a list of `Action`s.

```
type Action
| None
| For of Purpose
| Until of DateTime * Action
| After of DateTime * Action
| Choice of Action list
```

The type `Decision` signifies a binary choice of `Allow` or `Disallow` and is used to build the underlying form of the `Contract` datatype.

```
type Decision = Allow | Disallow
```

A `Code` is used to explicitly identify an item of data that is being queried against and is designed to hold a value translatable into an underlying clinical coding system. At present the type `Code` simply holds a string to be matched against. A more complex type definition of a clinical code is possible but was not deemed necessary for this prototype implementation.

```
type Code = CodeValue of string
```

On a basic level a `Contract` associates an `Action`, a list of `Code`s or both with a `Decision`. This captures the essence of the `Contract`, which is to allow or disallow given actions or requests for types of clinical code. `Contract`s which specify multiple `Action`s are grouped using the `Or` constructor and finally a default decision is specified using the `Other` constructor, although in practice and in the absence of a specific value this will be presumed to be `Disallow`.

```
type Contract =
  ForAction of Action * Decision
| ForCodes of Code list * Decision
| ForActionCodes of Action * Code list * Decision
| Or of Contract * Contract
| Otherwise of Decision
```

A `Request` mirrors in part the form of a contract, but pertains to only a single request with the corresponding `Or` constructor missing.

```
type Request =
  RequestFor of Action
| RequestForCodes of Code list
| RequestForActionCodes of Action * Code list
```

Finally a `Warrant`, produced as the result of a Request being matched against a Contract specifies that the holder of the `Warrant` be allowed to perform a given Action against the data held by the holder of the corresponding Contract. The Warrants constructor allows for a list of individual Warrants to be held in the data type.

```
type Warrant =
  WarrantedFor of Action
| Warrants of Warrant list
```

The programmatic API for accessing the system essentially consists of a single function that takes a Contract and a Request and produces a corresponding Warrant.

```
type VerifyRequest = Contract -> Request -> Warrant
```

The semantics of the function implementation are such that if the `Action` of the request matches a corresponding Action of a contract with a given `Allow` decision type then a `Warrant` for that `Action` is returned. If the corresponding decision type is `Disallow` then a `Warrant` is produced but with an `Action` of `None`. Similarly the underlying function implementation matches requests for `Codes` in a similar way, but the method for deciding whether two different `Code` strings match can vary based on the underlying semantics of the given coding system.

3. Results

We demonstrated the applicability of the system described above by translating existing information governance agreements into this computable format. These contracts were taken from real-world examples of information governance agreements placed on users of data from the Salford Integrated Record [8]. In the following we give a description in plain language of what the governance contracts specified in terms of the allowable forms of request and then show the corresponding implementation of a Contract datatype for two such contracts.

Contract 1: Data agreements were already in place for access to data for research purposes. Additionally access to data for auditing purposes would be granted for a six-month period. The contact granted benchmarking access to users for this period and then reverted back to the original agreement for ongoing research use. The corresponding formally typed definition is given below.

```
let contract1 = Or (ForAction (Until ("03-01-2016")
                                 (For Benchmarking)))
                   (ForAction (For Research)) Allow
```

Contract 2. A data agreement was put in place allowing research access to data pertaining to Diabetes or for Asthma. The underlying coding system used was the Read code system. The corresponding formally typed definition is given below.

```
let contract2 = ForActionCodes (For Research) ["CD10.."; 
"H33.."] Allow
```

The system was used to translate a series of such contracts into the underlying formal specification, demonstrating the general feasibility of the system for such use.

4. Discussion

The exemplar results outlined above demonstrate the feasibility of the system in terms of its ability to translate existing information governance agreements into a computational format. This ability removes the existing 'analogue gap' between electronic health records existing in the digital domain and information governance contracts that lie outside that domain and yet are used to enforce the use and control of digital health records. As it stands the system is in proof-of-concept stage, and whilst we have demonstrated its feasibility through the translation of existing real world agreements, fully integrating it into an existing system would serve to fully prove its suitability in the real world. Given that information governance has at its core an underlying legal foundation, there will likely be barriers to adoption in terms of a requirement for computational contracts as described here to be formally verified to meet such legal requirements. This, combined with the natural reticence that exists towards the adoption of novel methods to be applied to a field as sensitive as personal medical data may make moving forwards in the real world difficult.

Information governance is a vital aspect in the use of electronic healthcare data. Without strict information governance procedures in place the use of electronic health-care data, particularly for research purposes, will remain limited with respect to its potential use. By moving the specification of IG contracts into a computable form we have taken a step towards demonstrating the potential for formally enhancing the trust that can be placed in such systems. This is the first step in integrating the currently informal (in a computational sense) specification of IG procedures into the digital realm within which patient records already reside. Moving forwards we see this as a key future direction into the full digitisation of the healthcare domain, which we see as a crucial future direction of the field.

References

[1] House of Commons Health Committee, "The electronic patient record." http://www.parliament. the-stationery-office.co.uk/pa/cm200607/cmselect/cmhealth/422/422.pdf, July 2007. [Online; accessed 7-Nov-2016].

[2] A. Cornwall, "Electronic health records: An international perspective," *Health Issues*, no. 73, 2002.

[3] T. B. Murdoch and A. S. Detsky, "The inevitable application of big data to health care," *JAMA*, vol. 309, no. 13, pp. 1351–1352, 2013.

[4] R. Lewis, J. Smith, and A. Harrison, "From quasi-market to market in the National Health Service in England: what does this mean for the purchasing of health services?," *Journal of health services research & policy*, vol. 14, pp. 44–51, 2009.

[5] A. Geissbuhler, C. Safran, I. Buchan, R. Bellazzi, S. Labkoff, K. Eilenberg, A. Leese, C. Richardson, J. Mantas, P. Murray, *et al.*, "Trustworthy reuse of health data: a transnational perspective," *International journal of medical informatics*, vol. 82, no. 1, pp. 1–9, 2013.

[6] B. M. Knoppers and R. Chadwick, "The ethics weathervane," *BMC medical ethics*, vol. 16, no. 1, p. 58, 2015.

[7] M. J. Taylor and N. Taylor, "Health research access to personal confidential data in England and Wales: assessing any gap in public attitude between preferable and acceptable models of consent," *Life sciences, society and policy*, vol. 10, no. 1, p. 1, 2014.

[8] NHS Salford, "Salford integrated record. sharing patient information locally." http://www.salfordccg.nhs.uk/documents/Publications/SIRA5Booklet.pdf, July 2016. [Online; accessed 7-Nov-2016].