

On Computable Expressions of Policies for Digital Health: Use for Privacy Consent

Zoran Milosevic ^{a,b,1}

^a Deontik, Australia

^b Institute for Integrated and Intelligent Systems, Griffith University, Australia

Abstract. This paper proposes a formal model for expressing policies in digital health. The aim is to support computable expressions of legislative, regulative and organizational policies. The model is grounded in the semantics of deontic logic [1] and in modelling concepts for expressing accountability, specified in the new RM-ODP Enterprise Language standard [2]. An example of privacy consent based on the FHIR consent resource [3] is used to explain the use of these modelling concepts. The example involves multiple stakeholders and illustrates the complexity associated with the use of machine learning and artificial intelligence systems as part of healthcare delivery governed by informed consent policies.

Keywords. policy; privacy; consent; artificial intelligence; ethics; FHIR

1. Introduction

Digital health ecosystem is undergoing significant transformation at present, enabled by new technologies such as mobile devices, cloud platforms, new generation of machine learning (ML), artificial intelligence (AI), clinical decision support (CDS) systems, genomics solutions and so on. The ecosystem is also maturing in terms of the solutions available and new interoperability standards, in particular HL7 Fast Health Interoperability Resources (FHIR) [4] and SNOMED-CT [5]. There are also new proposals related to the clinical and administrative workflows [7][8].

These solutions increasingly require addressing cross-organisational and cross-jurisdictional challenges, the central theme of which is ability to clearly specify healthcare and information policies. This is the topic which is broader to that of information security and needs to accommodate expression of such policies as they constrain behaviour of various stakeholders, individually or as part of their interactions with other stakeholders. Many such challenges arise in the context of future intelligent health visions, while unlocking the innovation potential from health data. In particular, the use of AI in the healthcare context is already raising a series of important societal and ethical questions which we will need to address now, to ensure that intelligent health can deliver on its promise, respect existing norms and more importantly, helping us develop norms for some new issues that are starting to emerge [6].

This paper proposes the use of recently published RM-ODP Enterprise Language Standard (ODP-EL) [2] as a basis for providing computable expressions of policies to facilitate clear statements of constraints on parties' behaviour – including the dynamics of delegation in the delivery of healthcare, from a clinician to clinician, but also capturing responsibility of other parties, e.g. the providers of AI and CDS systems.

¹ E-mail: zoran@deontik.com.

Computable expressions of policies are also needed to guide dynamic clinical workflow capability, that would gather and use patient-related information during clinicians' observational and cognitive activities. Such a capability can be regarded as an actor in care delivery, governed by a set of policies and acting like an intelligent agent, sharing collective awareness about ongoing activities – representing the *context* of the care delivery. Context captures many situational factors surrounding the delivery of care for a specific patient, obtained through tests or clinician's observations, and managed using their cognitive models and evidence-based models of care. The context also includes jurisdictional and organisational policies that guide care delivery reflecting patients consent preferences – both from the operational and research aspects of the care.

Next section presents key concepts of the formal policy model based on the ODP-EL standard, as a UML meta-model. Section 2 illustrates use of these concepts to model privacy consent policies. Section 3 provides discussion and areas of future work.

2. Generic Policy Meta-Model

2.1. Policy Context

The central part in defining computable healthcare policies is the specification of constraints on the *actions* of the parties who participate in interactions. These constraints are prescribed by legislative, regulative or organizational authorities - defining applicable laws and rules for resources, data or interactions under question, i.e. policy context. For that purpose, the precise semantics of the ODP-EL concept of *community* can be used to describe the organizational or social environment for the participants. A community contract is defined in terms of community roles, their interactions and policy constraints that apply to the roles [9]. A community role can be fulfilled by an enterprise object which can be an IT system, a party (which models a natural person or legal entity), or another community, making it possible to model hierarchical policy contexts.

2.2. Deontic Constraints

There are three fundamental types of policy constraints in any normative system:

An *obligation* is a prescription that a particular behaviour is required. An obligation is fulfilled by the occurrence of the prescribed behaviour. A *permission* is a prescription that a particular behaviour is allowed to occur. A permission is equivalent to there being no obligation for the behaviour not to occur. A *prohibition* is a prescription that a particular behaviour must not occur. A prohibition is equivalent to there being an obligation for the behaviour not to occur.

These definitions have been the subject of standard deontic logic [1], but their application in designing enterprise systems requires explicit association with the agent to which these constraints apply. This is also needed to accommodate an agent's goal-seeking behaviour, which may result in their willingness to violate the policies with the expected benefit of potential future reward from doing so [2]. The way that deontic constraints are associated with the agents (i.e. active enterprise objects in ODP speak) is through *deontic tokens*. These are enterprise objects which encapsulate deontic constraint assertions. The holding of the deontic tokens by active enterprise objects constrains their behaviour. This modelling approach provides a pragmatic means for manipulating deontic tokens, for example, passing them between parties to model delegations, and activation or de-activation of policies that apply to the active enterprise objects in the context of their enterprise interactions. There are three types of deontic tokens: *burden*,

representing an obligation, *permit*, representing permission and *embargo*, representing prohibition. In the case of a burden, an active enterprise object holding the burden must attempt to discharge it either directly by performing the specified behaviour or indirectly by engaging some other object to take possession of the burden and perform the specified behaviour. In the case of permit, an active enterprise object holding the permit is able to perform some specified piece of behaviour, while in the case of embargo, the object holding the embargo is inhibited from performing the behaviour (see Figure 1).

In order to support the passing of deontic tokens among objects such as patient giving permit to a researcher to access their health record, the concept of a *speech act* is introduced. This is a special kind of action that is used to modify the set of tokens held by an active enterprise object. The name was chosen by analogy to the linguistic concept of speech act, which refers to something expressed by an individual that not only presents information but performs an action as well. Thus, a speech act changes the state of the world in terms of the association of deontic tokens with active enterprise objects. This modelling feature fits well with the nature of AI-enabled digital health applications, because it can support traceability of obligations of parties (clinicians and AI system creators), according to their broader responsibilities derived from ethical, social or legal norms, as further refined through the accountability concepts, described next.

2.3. Accountability Concepts

Party is as an enterprise object which models a natural person, or any other entity considered to have some of the rights, powers and duties of natural person, for example, a company. ODP-EL introduces two other concepts which are useful to describe many forms of delegation in enterprise systems: *Principal* is a party that has *delegated* something (e.g. authorization or provision of service) to another, and *Agent* is an active enterprise object that has been delegated something (e.g. authorization, responsibility of provision of service) by, and acts for, a party (e.g. in exercising the authorization, carrying out responsibility). *Delegation* is an action that assigns something (e.g. authorization, responsibility of provision of service) to another object, such as the act of referral. It is through this mechanism that deontic tokens can be passed across different active enterprise objects, with one example being a delegation from principal to agent.

There are several other action types to capture important business events in any organizational system, and reflect the dynamics of communication amongst parties, and broadly, active enterprise objects [2]. *Commitment* is as an action resulting in an obligation by one or more participants in the act to comply with a rule or perform a contract. This effectively means that they will be assigned a burden. Examples include commitments by clinicians to deliver safe, reliable and effective healthcare to patients. *Declaration* is defined as an action by which an object makes facts known in its environment and establishes a new state of affairs in its environment. This can, for example, be performed by an AI system (or a party managing it), for example, informing the interested parties about the result of some analysis.

Prescription is an action that establishes a rule. Prescriptions provide a flexible and powerful mechanism for changing the system's business rules at runtime, enabling dynamic adaptation to respond to business changes and new needs. This ability is important in any digital health system, to establish the applicability of new policies reflecting new legislations for example, or after the adoption of new recommendations from AI system components.

Authorization is as an action indicating that a particular behaviour shall not be prevented. Unlike a permission, an authorization is an empowerment. In terms of deontic tokens, the enterprise object that has performed authorization will issue a required permit

and will itself undertake a burden describing its obligation to facilitate the behaviour. For example, the authorization for the consumer to challenge AI decisions is giving them permit to do so by the AI system (or its creator/manager) who has the burden to do so.

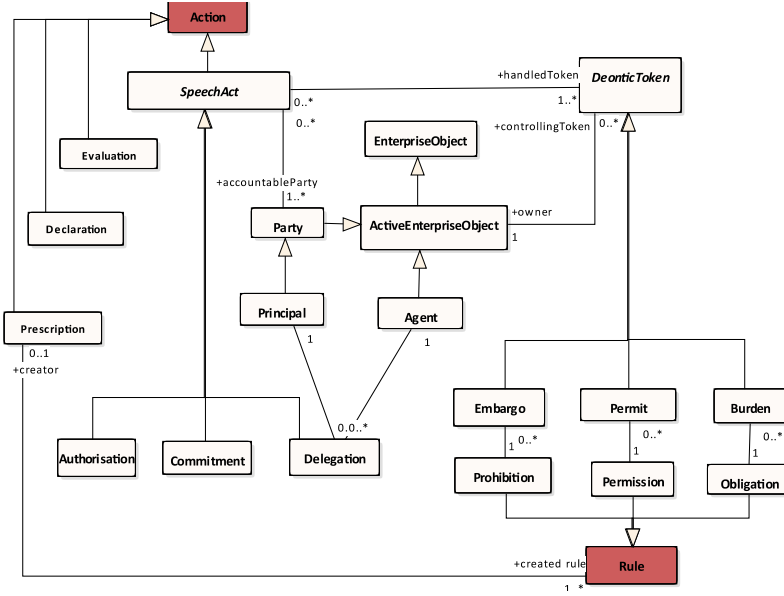


Figure 1. Key deontic and accountability concepts.

3. Privacy Consent Model

3.1. Policy Context

Recall that the policy context can be modelled using a community concept and thus a policy consent community specifies the following community role types (see Figure 2):

- *Grantor*, to be fulfilled by any individual giving consent under a set of permission rules, such as being of legal age.
- *Grantee*, to be fulfilled by professionals with the required credentials, which can be either a *Clinician*, with permission to access Grantors individual health information for care purposes (covered by the patients consent for primary care, e.g. access to all of the patient information in an emergency situation, with certain constraints, such as time period from the emergency event) or a *Researcher*, with permission to access Grantors de-identified health data for research purposes and obligation not to perform re-identification of patient data, as prescribed by National Data Protection Authority.
- *Consent Authority*, a trusted party responsible for storing individuals' consent agreements and overseeing the consent agreement rules; its function can also be to facilitate ethics approvals to govern the secondary use of data.
- *Research Broker*, a commercial entity authorized to search patient health data and consent data to identify patients suitable for research projects. The Broker is responsible to ensure that patient preferences are enforced. It is accountable to the Consent Authority and the National Data Protection Authority.

- *National Data Protection Authority*, responsible for defining and enforcing data protection policies, as legislated.
- *Automated Decision-Maker*, performing analytics, recommendations and in some cases, active decision-making; this role guides and augments activities of clinicians, researchers and other stakeholders, such as population health experts; this role can be fulfilled by clinical decision support systems or AI systems.

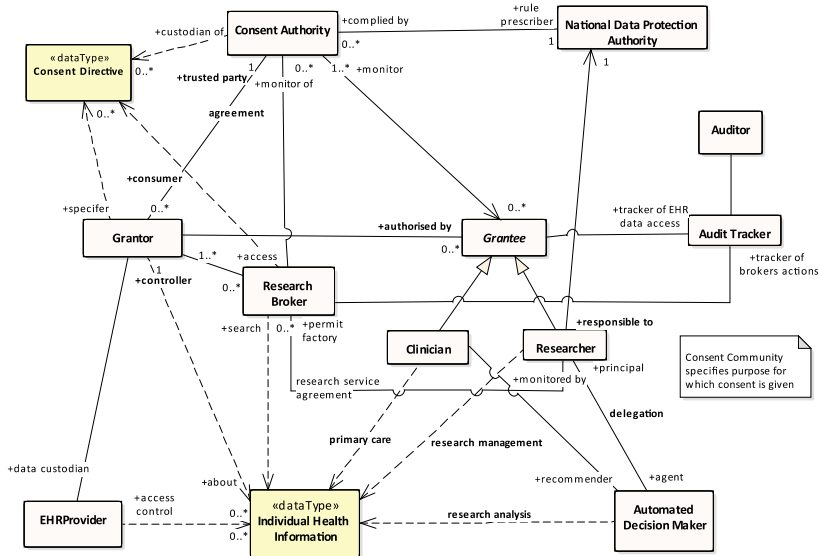


Figure 2. Privacy consent community.

3.2. Deontic Constraints

The privacy consent community defines a number of deontic constraints, such as:

- *Permission* of the Grantor given to the Consent Authority to store consent agreements, e.g. valid for a specified time period defined by the Grantor.
- *Permission* of the Grantor to the Broker to search patients' data and if it satisfies researcher criteria include a link to this data in a data set for the researcher.
- *Obligation* on the Audit Tracker to log data access by the Grantee reliably and on-time and provide access to the audit trail by the Auditor; the tracker may also have an obligation to log actions of Research Broker for forensic purpose.
- *Authorization* of the Grantor to the Grantee to access the Grantor's health information, as per the following actions: a) Grantor issues permit to the Broker for searching their data to establish whether they satisfy research question criteria, b) Research Broker issues a research permit to the Researcher which includes a list of Grantors that provided consent to access their de-identified health data and whose data satisfy the research question, c) EHR provider provides access permit to the Researcher to access health records of specific patients, provided researcher has credentials requested by the EHR provider.

3.3. Accountability Concepts

Authorization is modelled using a combination of permit and burden deontic tokens. For example, authorization of the Grantor to the Broker involves both the permit being passed from the Grantor to the Broker to search its record but also places an obligation on the Grantor itself, through the corresponding burden, to ensure that access to its record is ultimately enabled. This authorization action is also a speech act because it changes the deontic state of both the Grantor and Grantee. The effect of this speech act is that the existing grantor's permit to the Broker to search its healthcare data is passed on to the Grantee. In this example, we assume that the consent directive gives permission to the researcher to access the Grantors health data but prohibits access to the Grantor's mental health data (if it exists). The use of speech acts and deontic tokens is a convenient means for describing the dynamics of deontic constraints and passing of tokens, including to the parties with ultimate legal responsibility. Many data protection rules defined by a National Data Protection Authority set accountability and legal responsibility expectations for actions of researchers involved in using grantor's data. These data protection rules were established through *prescription* actions), performed by the National Data Protection Authority, which essentially establishes obligations and permissions for all the parties involved in accessing patient data.

4. Discussion and Future Work

This paper presents an approach to a computable expression of healthcare policies. This is a difficult problem, but we believe the use of precise modelling framework provided by the ODP-EL standard offers a promising solution path, supported by the use of contemporary software modelling tools. We illustrated this through the example of privacy consent, which demonstrated the expressiveness of the approach, in spite of the limited space available in this paper. In future, we plan to consider applying this foundational policy model to the recently proposed dynamic consent model [10] and consider personalised and fine-grained controls over access to individual information. We also plan to investigate in detail patient's specific consent related to the purpose for which clinicians may use AI systems [11] as part of their care delivery for patients, as well as for providing constraints over analytics applications [13]. Finally, we plan to use this policy model to further investigate ethics and legal challenges associated with responsibility of using AI in digital health as initially proposed in [12].

References

- [1] G.H. von Wright, Deontic Logic, *Mind*, Vol 60, pp. 1-15, 1951.
- [2] *ISO/IEC 15414*. 2015. Information technology: Open distributed processing, Reference model, Enterprise Language, 3rd ed.
- [3] *HL7 FHIR Consent Resource*, Release 4, <https://www.hl7.org/fhir/consent.html>, (Accessed 7 Aug. 2020).
- [4] *HL7 FHIR* <https://www.hl7.org/fhir/index.html> (Accessed 7 Aug. 2020).
- [5] *SNOMED-CT* <http://www.snomed.org/snomed-ct/why-snomed-ct> (Accessed 7 Aug. 2020).
- [6] *Microsoft*, Healthcare, artificial intelligence, data and ethics - A 2030 vision, Dec 2018.
- [7] *Object Management Group (OMG)*, Business Process Management for Healthcare (BPM+ Health), <https://www.bpm-plus.org>, (Accessed 31 Dec. 2019).
- [8] *HL7 FHIR Workflow Definition*, <https://www.hl7.org/fhir/workflow.html> (Accessed 7 Aug. 2020).
- [9] Linington, P., Milosevic, Z. Tanaka, A. & Vallecillo, A., Building Enterprise Systems with ODP, An Introduction to Open Distributed Processing. *Chapman Hall/CRC Press*, 2011.
- [10] Dynamic consent, https://en.wikipedia.org/wiki/Dynamic_consent.

- [11] An invisible hand: Patients aren't being told about the AI systems advising their care <https://www.statnews.com/2020/07/15/artificial-intelligence-patient-consent-hospitals/>
- [12] Milosevic, Z., Ethics in Digital Health: a deontic accountability framework, *Proc. IEEE EDOC'19 Conf.*, Paris, 2019.
- [13] A Berry, Z Milosevic, Real-time analytics for legacy data streams in health: monitoring health data quality, *Proc. IEEE EDOC'13 Conf.*, Vancouver, Canada, 2013.