Digital policies for patient consents: the thorny (and general) technical challenges

Arnon Rosenthal, PhD.
The MITRE Corporation[1]
Extended Abstract

1. Introduction

To protect patient privacy, release of patient records must be in accordance with patient consents to share clinical data, either explicitly or from a government default. Explicit consents allow a patient to customize the balance between confidentiality versus sharing. However, the system of request-specific paper consent forms already acts as sand in the gears of data exchange [GSK], and will become far more bothersome as data sharing becomes nationwide, and progress on electronic health records gradually automates data extraction and transmission.

Our project is architecting and prototyping key elements of a system to elicit and manage consents. All of a patient's consent rules are to be managed in one place, editable over the web, and accessible by authorized record holders; the user interface will help the patient manage their consent rules. Then when an information request is received, it will put the appropriate set of rules on the record holder's screen, in human-readable and automation-friendly forms.

For robustness as systems evolve, our approach works with today's largely manual systems (espically at small health providers), where consent rules are enforced manually. As clinicians automate, tighter integration will provide better efficiency. It is intended to scale nationwide, in an environment with competing consent providers. A fuller discussion of the operational and policy issues to be confronted appears in [SM], also submitted to this workshop.

In the course of our work, we abstracted three fundamental technical challenges that arose for patient consents, and also are relevant in other digital policy arenas. This paper presents these problem formulations and our progress on them.

First, consent systems must mix preferences from patients and several layers of government, some expressed as defaults and others as legal mandates. Since interests may conflict, we need a way to easily and safely administer such mixtures, and to give the record holder something unambiguous to execute.

Second, patients may express releasability conditions (e.g., "contains no information relevant to substance abuse") that cannot be enforced 100% accurately. (Software cannot understand all the meaning and implications in diverse medical records, especially natural language; medical records staff are also fallible).

Third, record holders must see the consent rules in order to enforce them, but the rules themselves can contain restricted medical information.

Sections 2-4 below address these challenges in turn. Each is abstracted as a general technical problem of digital policy.

2. Specifying and merging patient and governments' rules

There are many stakeholders who may contribute rules to be applied to a data request, notably the patient, his/her home state, states where records reside, and the federal government.[1] Some rules say Allow release, others Deny. Ambiguities need to be resolved without pairwise negotiation among parties (1225 pairs of states, each with multiple agencies).

As a further complication, government opinions are a mix of absolute mandates (release certain infectious disease reports to public health departments), strong defaults (release substance abuse information only upon explicit permission), and weak defaults that apply in the absence of another opinion (Opt-Out: release all information to treating clinicians unless otherwise specified). Meanwhile, patients' opinions would be captured via a mix of explicitly-written rules, user interface shortcuts (e.g., "use medium privacy"), and clicking on legal boilerplate, each of which may deserve a different strength when others conflict.

[1] The patient's clinicians are also influential, but probably will not contribute explicit rules. They exert influence by designing consent forms, and warning the patient of the medical dangers of withholding drug prescriptions, allergies, etc.

Our technical challenge is to provide a framework for managing these issues, enabling regulators to make the legal and medical policy choices. We have drafted a rule formalism and operations concept that accommodates all these concerns, usually without burdening the patient. The major ingredients are:

- Both ALLOW and DENY rules, each with a strength attached. The combined ruleset will need to give an unambiguous answer.
- A predefined vocabulary of strengths, organized into three tiers (Explicit, Aware, Unaware) corresponding to degrees of patient involvement.
- A *jurisdiction wizard* that determines and prioritizes the stakeholders whose rules should be used in processing this request.[2]

Rules with weights are simple technically, and have been proposed many times. e.g., in [JSSS]. They have found limited use, probably because administration is difficult for one person, and even harder with autonomous stakeholders. The three fundamental difficulties appear to be ill-defined strength scales, stakeholders who game the system, and keeping it simple for the unskilled.

Consent seems to fall into a sweet spot, less vulnerable to these difficulties. First, the various governmental regulators already "play nice" by declaring some rules to be defaults that patients can override. Several federal defaults defer to states. Regulators seem likely to express genuine strength of opinion, not to game the system to "win" over other stakeholder states.

Second, government rule-writers can calibrate rule strength relative to natural tiers of strength for patient opinion, based on level of awareness and explicitness. For example, "All information is releasable to my primary physician" is more explicit than "Medium Privacy" or "I accept XYZ hospital boilerplate". Governments may then adjust based on their willingness to defer to others.

Third, patients can be insulated from most complexities. They almost never need to specify strengths– a user interface can report "awareness and explicitness strength" based on how the rule was obtained. It remains challenging to explain the behavior of a rule system to technologically naïve users, but multiple stakeholders do not greatly worsen it.

Initially, we had just two stakeholders specifying privacy rules, Patient and Government. Later, we elaborated to have rules from multiple government entities. We also found a need for workaround constructs to allow government experts to adjust behavior in important special cases, notably to nullify a lower agency's Deny rule without expressing a countervailing Allow, or to let the federal government prevent states from overriding a patient choice in a particular case.

We tested our system's flexibility and ease of use by writing rules that handled many plausible stakeholder interactions. We were able to make nearly all the distinctions we wanted, and usually succeeded in keeping simple things simple, and difficult ones possible. While a few cases could not be handled straightforwardly, the ones we did make easy seem a large advance over current practice.

The problem of stakeholder interactions is widespread. It is an open question how well our approach (strength tiers + adjustments) would work in domains without the favorable administration characteristics listed above.

3. Don't just seek perfection – manage imperfection

Privacy advocates and medical informatics experts [PCAST] urge that patients be given fine grained control over the contents of each data release. For example, a patient should be able to exclude mental health information from routine releases, and to absolutely exclude the clinic where an estranged spouse works, even from emergency "break the glass" access.

Unfortunately, neither medical records personnel nor technologists can guarantee to enforce these restrictions with high accuracy. Difficulties include determining the meaning of text, determining what a clinician might infer from prescriptions and symptoms, inability to interpret records imported from other providers, and clinics doing business under multiple names. A

[2] Since jurisdiction rules may change, the architecture makes it a separate module. We examined operational tradeoffs based on what metadata the decision is allowed to depend on.

further complication is that patients differ in what they consider very sensitive. Even if clinicians had the time to tag many topics, some patients would want different ones, or more specific ones. (For example, a patient might want to redact only abortions, but share other information tagged "reproductive health").

In fact, no conceivable technical progress and no feasible manual review regime will guarantee perfection. Thus, we must manage an imperfect world. Below, we first examine how record holders and patients may react, and then propose "consent for release with precautions".

Responses by record holders and patients. When they know they cannot enforce perfectly, record holders will respond in different ways, making the effect of the patient's rule quite unpredictable. For example, consider the rule "Allow access by doctors at Clinic C except *release mental health data only to Dr. Freud*".

Some record holders may simply refuse to share any data under such restrictions, fearing blame if some mental health data slips through. Others will share only data their software can be rather sure of, such as data collected locally, under a protocol that tagged all information relevant to mental health. Older data and data from outside consultants might be withheld. Patients thus do not receive desired data sharing.

Still others, believing that sharing is critical to care, will apply reasonable precautions and then send the data. Such a record holder is ill served today, having no way to tell which patients feel extreme needs for confidentiality (e.g., due to a child-custody lawsuit or facing legal jeopardy).

Patients' options also suffer from such blunt instruments. If they *strongly* want data to be shared, even by cautious record holders, they can omit the specific restriction (here, on mental health). But now they are not even asking record holders to *try* to filter. Alternatively, they can omit the Allow rule that permitted restricted sharing, thus blocking sharing entirely. These are both deeply unsatisfactory.

Ameliorating the problem. At a minimum, patients should be able to express "when sharing, please try to block the following" without imposing a legal requirement that filtering succeed. Conversely, record holders should have incentives to deploy filters and apply practices that are helpful, albeit imperfect. Some record holders (especially IT services like Google Health) might treat ability to apply stronger protections as a competitive advantage.

To mitigate these difficulties, we then propose 'consent to release with a specified level of precautions". To allow uncertainty to be managed intelligently, we will need a legal framework, common vocabularies, consent constructs, and user interface help. Specifically, we propose that:

- When capturing a rule with a difficult restriction, the consent UI should also ask the patient to state their degree of concern for perfect enforcement. The system can map it to a set of precautions.[3]
- Government could encourage creation of standard vocabularies to describe individual precautions, and also packages of precautions to be considered level 1, 2, etc. It should also encourage measurement and certification of techniques' effectiveness.
- Record holders who implement these precautions (or preferably, better ones) would be considered legally compliant.

One can imagine many feasible, low cost precautions. The first level might be to instruct staff "while carrying out normal procedure, try to avoid this topic". Beyond this, automated filters might look for "dirty words" such as "delusions" and omit any treatment note that includes it. One can infer much from diagnostic codes. A knowledge base might say "a prescription for drug X suggests sensitive diagnosis Z". A good practice might be to send only 1 year's data, unless the requestor explicitly asks for more.

Several fundamental consent conundrums were out of scope for our "imperfection" investigation. The tradeoff between privacy and risk of inhibiting sharing arises even if filtering is perfect. The same tradeoff applies to notifying recipients that data has been redacted.

4. Limiting release of sensitive Consent documents

---

[3] The system should also elicit whether the patient is trying to hide the detailed records or the very existence of a topic. In the latter case, more effective precautions are required, since a single record's release reveals existence.

The statement of a consent rule can itself contain sensitive information. For example, one might not want to send the following rules to your dietician or nursing home:

- Do not forward any of my records that originated at ABC Rehab or XYZ prison.
- Do not forward any records that mention my contraception.

The Consent rules can be seen as yet another set of protected health data, but they need special treatment, for at least two reasons:

- A record holder that does not receive restrictions cannot enforce them.
- Requests for consent information will not receive human review. Automation is needed to meet speed and cost requirements.

There are many ways to take wedges out of the problem:

- Some requests are not allowed, based on the request message's recipient, purpose, etc., without checking medical records. Record holders need not be sent the applicable ruleset.
- One could replace the specific restriction by a broader, less informative one, e.g., Reproductive Health for contraception, or to withhold all information across a conventional list of sensitive categories.
- The consent system may know that certain record holders are highly unlikely to hold information covered by the restriction, e.g., that a dietician says that she never has abortion data. Other systems might advertise that they never release certain topics without very explicit permission. Such system need not be sent Deny rules for these categories.
- Sensitive institutions such as ABC Rehab might tag all their data "do not disseminate further". (A standard construct is needed).

References:

[GSK] N. Genes, J. Shapiro, G. Kuperman "Health Information Exchange Consent Policy Influences Emergency Department Patient Data Accessibility", *Proceedings of AMIA Symposium*, 2010

[JSSS] S. Jajodia et. al., "Flexible support for multiple access control policies", *ACM Trans. Database Systems, June 2001*.

[PCAST] President's Council of Advisors on Science and Technology, "Realizing the full potential of health information technology to improve healthcare", December, 2010 http://www.whitehouse.gov/sites/default/files/microsites/ostp/pcast-health-it-report.pdf

[SM] J. Stanford, P. Mork, "Patient Health Data Consent Management on a Nationwide Scale: What Needs To Be Done?" submitted to *HealthSec '11*

---

[1] MITRE is a nonprofit corporation that operates Federally Funded Research and Development Centers.

The author wishes to acknowledge the help of Jean Stanford, Peter Mork, and Linsey O'Brien.