

Reasoning about the Appropriate Use of Private Data through Computational Workflows

Yolanda Gil and Christian Fritz

University of Southern California, Information Sciences Institute
Marina del Rey, CA, USA
{gil, fritz}@isi.edu

Abstract

While there is a plethora of mechanisms to ensure lawful access to privacy-protected data, additional research is required in order to reassure individuals that their personal data is being *used* for the *purpose* that they consented to. This is particularly important in the context of new data mining approaches, as used, for instance, in biomedical research and commercial data mining. We argue for the use of computational workflows to ensure and enforce appropriate use of sensitive personal data. Computational workflows describe in a declarative manner the data processing steps and the expected results of complex data analysis processes such as data mining (Gil et al. 2007b; Taylor et al. 2006). We see workflows as an artifact that captures, among other things, how data is being used and for what purpose. Existing frameworks for computational workflows need to be extended to incorporate privacy policies that can govern the use of data.

Introduction

Individual privacy is an unquestionable right whose preservation and enforcement present great challenges. Table 1 summarizes the eight principles put forward by the widely-cited Fair Information Practices of the Organization for Economic Cooperation and Development (OECD) (OECD 1980), which are the basis for privacy laws in many countries including the US, the UK, and Germany as well as the EU. Some of these principles (collection limitation, data quality, and individual participation) are concerned with data acquisition and management. Most of the principles, however, concern the *use* of data, and raise fundamental research questions regarding the technologies and architectural mechanisms to implement them, such as: *Use limitation*: What technologies can limit the use of sensitive data to those uses allowed by privacy policies? *Purpose specification*: How can the purpose of the collection of data be disclosed in an intelligible manner to individuals? What technologies can ensure that the use of the data is consistent with this purpose? *Openness principle*: What technologies could reassure individuals about the implementation of privacy policies restricting the use of information? *Accountability*: What technologies could be applied to enforce privacy policies concerning the use of sensitive information?

There has been prior work on addressing privacy concerns in that regards the use of the data. Some research to date has focused on access control mechanisms and

standard policy languages (XACML 2005; SAML 2005; P3P 2002). Other research has focused on protection of individual information through anonymity, e.g., (Sweeney 2002; Jiang and Clifton. 2006) de-identification (Uzuner, Luo, and Szolovits 2007), and abstraction (Kargupta et al. 2005). Recent research has led to privacy-preserving data mining algorithms that protect sensitive data, e.g., (Agrawal and Srikant 2000; Lindell and Pinkas 2002; Jiang and Clifton. 2006). Other research provides techniques for managing sensitive data collections, e.g., (Esponda et al. 2006).

We propose to use computational workflows to answer these questions and as a means to express and monitor the appropriate use of data. Computational workflows are increasingly used in a variety of applications, notably in e-Science, to represent and manage complex computations with distributed data sources and execution resources (Gil et al. 2007b; Taylor et al. 2006). Each step in the workflow is a data analysis or data integration step. These steps are linked according to their dataflow. Workflows can describe the individual and overall computations used to process distributed data, the purpose of the computations, the intermediate and final results, and the original information sources.

We focus on workflow systems as data analysis frameworks that have many choices regarding data sources to use and types of analyses to conduct. Left to its own devices, such a system could make any use (i.e., run any process) with the data it has access to. However, as a user or the system faces those choices, it should be driven to choose algorithms and data sources that satisfy privacy policies. Privacy policies could constrain the overall system behavior and prevent it from making choices that are not in violation of those policies. We are not suggesting that any organization using a workflow system to reassure others of their use of sensitive data would not be able to maliciously break privacy laws or consent agreements. Rather, we aim to provide for the possibility of accountability.

Although most emphasis on privacy is on the controlled access of data, we see tremendous benefits to having computational mechanisms that ensure the lawful use of data once accessed. First and foremost, individual privacy would be better protected if the use of the data was controlled and auditable. Second, we hypothesize that more individuals would be willing to allow access to personal data if such enforcement mechanisms were in place. Even though many are reluctant to allow the use of their personal data, many

Principle	Description
Collection limitation	The collection of personal information should be limited, should be obtained by lawful and fair means, and, where appropriate, with the knowledge or consent of the individual.
Data quality	Personal information should be relevant to the purpose for which it is collected, and should be accurate, complete, and current as needed for that purpose.
Purpose specification	The purposes for the collection of personal information should be disclosed before collection and upon any change to that purpose, and its use should be limited to the purposes and compatible purposes.
Use limitation	Personal information should not be disclosed or otherwise used for other than a specified purpose without consent of the individual or legal authority.
Security safeguards	Personal information should be protected with reasonable security safeguards against risks such as loss or unauthorized access, destruction, use, modification, or disclosure.
Openness	The public should be informed about privacy policies and practices, and individuals should have ready means of learning about the use of personal information.
Individual participation	Individuals should have the following rights: to know about the collection of personal information, to access that information, to request correction, and to challenge the denial of those rights.
Accountability	Individuals controlling the collection and use of personal information should be accountable for taking steps to ensure the implementation of these principles.

Table 1: The Eight Principles for Fair Information Practices (OECD 1980), many concerning the use of sensitive data.

would accept a loss of their privacy for some greater good such as advancing medical research or contributing to national security, as long as their choice to release personal data for those purposes will not result in uses that those individuals would consider non-desirable.

In the next section we further motivate the need for technology that helps protect privacy with respect to the use of sensitive information. Following that, we summarize the results of some of our preliminary work regarding the use of workflows systems to protect privacy. From that we then derive open questions and requirements.

Motivation

We provide more specific motivation for the need for privacy policies constraining the use of data from three major perspectives: e-science, data mining, and the open Web.

e-Science

The benefits of sharing clinical patient records across health care providers have long been recognized, together with the importance of protecting personal privacy (Dugas et al. 2002). Examples of uses of medical records include cancer (cabig.nci.nih.gov) and neuroscience (www.nbirn.net) research. Many advocate that patients should be in control of the collection and access of their own records, and be able to choose to keep some clinical data private even at the risk of getting inadequate treatment (Kohane and Altman 2005). In order for patients to allow sharing of their personal health data, there must be appropriate technologies available that capture and enforce their wishes. A variety of mechanisms are being investigated to ensure privacy in patient records including secure data storage, data access control, auditing mechanisms, and securing lines of communication (e.g., to safeguard privacy in the increasingly frequent use of e-mail exchanges between patients and their physicians), e.g., (Uzuner, Luo, and Szolovits 2007; Kohane and Altman 2005). At the same time, laws and policies for protecting and enforcing health information privacy

will need to be formulated in order to determine how those technologies need to be used to implement the law.

These mechanisms are important and necessary to control the access and release of data. However, they will not necessarily support the anticipated sophistication of patient's wishes over the fine-grained control over the uses of their clinical data. Kohane and Altman (2005) and Mandl, Szolovits, and Kohane (2001) have observed altruism in patients regarding their willingness to grant very restrictive kinds of data access permissions based on the specific study to be conducted. When it is for the advancement of medical knowledge and may, for example, help to cure cancer, many people are happy to make their clinical data available. However, ensuring that this use limitation is respected is becoming increasingly more complicated as the uses of clinical data become more and more distributed in nature and are not confined to the institution that collected or owns the data.

We argue that computational workflows enable the expression and enforcement of the kinds of policies that would give patients the very desirable control over the use of their clinical records, rather than access alone, and safeguard scientists from the unlawful use of private data, even when studies are distributed over many independent research labs.

Data Mining

Government and commercial use of data mining raises many issues concerning the appropriate use of private citizen information. In the area of government use of data, there is a vivid debate about trade-offs between providing security and the right to privacy. For example, Weitzner et al. (2008) present a scenario of violation of privacy laws where data collected for the purpose of airline passenger screening should not be used for the enforcement of other criminal laws. Other concerns arise from the acquisition of sensitive data by the government from the private sector, including personal communications (e.g., telephone calls, emails, etc.) and personal activities (e.g., location and context provided by mobile devices). In e-commerce,

privacy concerns result from limiting the release or resale of financial records and web activity. A particular area of concern is record linkage technology to cross-reference independent data sources and data mining to detect patterns and associate them with individuals (Sweeney 2004; Weitzner et al. 2008).

In order for the public to be reassured of the use of personal information by the government and the private sector, we need better mechanisms to demonstrate what data they are using, what data they are combining, and for what purpose in whatever data mining systems of routine use. There should also be technologies that enable them to prove that they have reached a conclusion or taken an action based on lawful means that respect individual privacy by showing provenance records of how results were obtained.

The Open Web

The Web opens additional concerns for privacy data. The wide availability of yellow pages and other directory information enable the re-identification of records that were previously anonymized (Weitzner 2007). In addition, protected or sensitive information may become available over the Web perhaps unintentionally.

Trust and security were always central to the vision of the Web (Berners-Lee et al. 2006) and these concerns receive renewed attention with the proliferation of social networks. People, it seems, are more than happy to post in public forums all sorts of personal opinions and details about their life, but are rightly upset when someone else uses personal information that they consider private. Individuals can control within a social networking site what they share and with whom. But as content is increasingly shared across networks and initiatives such as Linked Data (<http://linkeddata.org>) facilitate a diversity of applications based on this, new technologies are needed to carry out that control to external uses of the data.

Instead of advocating for overly restrictive privacy laws that would likely end up being counter to the thirst for sharing of the individuals they would be designed to protect, Weitzner et al. (2006) propose developing systems that are transparent and accountable regarding their use of sensitive data from individuals and therefore can demonstrate their compliance with existing privacy laws. Individuals must receive reassurance that their personal information that appears on the web will not be used for purposes that are unlawful or that would be considered an intrusion to their privacy. Moreover, if information about individuals that is available on the web is combined with other data sources (e.g., commercial, government) then undesirable intrusions on privacy may occur. Any organization should be held accountable for their use of personal data from the web. This requires a workable definition of lawful and unlawful use of sensitive information and we argue that computational workflows are one possible mechanism for this.

Privacy Protection through Computational Workflows

In this section we discuss preliminary work that enabled us to understand the issues involved in managing privacy through workflows. We used the Wings workflow system (Gil et al. 2007b; 2010) to represent and reason about data in the context of workflows and privacy.

Semantic Representation and Reasoning in the Wings Workflow System

Wings represents typical analysis methods as abstract, highly reusable workflow templates. Templates compactly express parallel processing of data collections even before the collections are selected. A key aspect of Wings is its workflow generation algorithm, which reasons about semantic properties of data and algorithms to generate valid workflows. Given a dataset to be processed, Wings reasons about the properties of that dataset, and reasons about how they satisfy the requirements and constraints of each step (or component) in the workflow template. It also uses predictive rules that express what parameter settings are most appropriate for the user's dataset. Detailed records of how new data products were generated by the system are captured in a provenance catalog. Wings generates workflow candidates and elaborates them until they specify the needed computations (but not where they will take place) and their dataflow dependencies. Wings then submits the workflows to the Pegasus workflow execution and mapping engine to assign run-time execution resources (Deelman et al. 2005). Pegasus performs a mapping of the workflow computations onto the resources based on the execution requirements of the codes and data, and adds steps to the workflow to carry out any data movements across locations. Pegasus then submits the workflow to the Condor DAGMan execution engine and monitors its execution, performing dynamic workflow remapping for some types of run-time failures.

We argue that the following key aspects of Wings are important for the development of the proposed privacy-aware workflow framework:

Semantic representations of attributes of datasets as well as models of workflow components, which the system uses to reason about workflows. This is important in order to reason about privacy-related attributes of the data, and to reason about the transformations that privacy-preserving algorithms perform on the data.

Distributed workflow execution, where data can be available in separate locations and the execution of workflow steps can also be distributed in a way that is controlled by the system. This is important in order to model privacy-preserving transformations that need to occur at the data source and what information is moved out of the location where the original sources are.

Provenance record keeping, where any aspect of the workflow creation and execution can be recorded for later examination to provide transparency over the system's operation. Provenance records allow us to expose what the

workflow system did to obtain any given result, thus supporting accountability.

Reasoning about Privacy Policies in Wings

We created a prototype of a workflow system that checks privacy policies for workflows based on Wings (Gil et al. 2007a; Cheung and Gil 2007). The workflows describe how data is used in terms of how it is analyzed and processed. To exemplify applications that could raise privacy concerns regarding use, we modeled data mining algorithms that could be used as workflow steps, called *components*, and created semantic representations of data and workflows that use those components. Both, components and data were described in OWL/RDF.

We first defined a *component catalog* that contained a range of data mining algorithms as well as privacy preservation techniques. The catalog was not meant to be exhaustive, but rather be representative of the kinds of algorithms that are relevant to reasoning about privacy. Data mining algorithms included clustering methods (e.g., k-means, Gaussian mixture models), manifold learning (e.g., GTM), and classification (e.g., SVM). Privacy preservation techniques were divided into two subclasses: per attribute and per dataset. The former had several subclasses including anonymization, perturbation, and encryption. The class of privacy preservation techniques per dataset included generalization algorithms such as k-anonymity. We also defined a *data ontology* with semantic representations of datasets, which essentially provided a meta-data vocabulary that we could use to reason about how datasets are transformed by the workflow components upon execution. Roughly, attributes of datasets had associated properties that expressed whether the attributes were protected by privacy preservation methods (e.g., whether they were anonymized). In addition, domain-specific ontologies were used to express the use that was authorized by the individuals when the data was collected. Using this data ontology, we populated a *data catalog* with initial datasets and specified meta-data attributes and values using the ontology. Finally, we defined workflows whose computational steps were elements of the component catalog and whose input datasets were elements of the data catalog. We defined rules that would represent reasonable constraints to address privacy protection. Each rule had a context that referred to the condition where the underlying policy was relevant, so that the policy applied only if this condition was satisfied, and a set of requirements that represented non-amendable conditions under which the use of data was required or not allowed.

Figure 1 shows two example workflows. Suppose that separate hospitals have data about a clinical trial that a researcher wants to analyze. Workflow A takes data from different locations into a separate site and performs a clustering analysis (using a Gaussian Mixture Model) after aggregating the data. Workflow B carries out anonymization, k-anonymity, and abstraction in each original site and then does the aggregation and clustering in a separate site. The second workflow respects a number of privacy policies that the first one does not. The following example rule asserts that datasets containing drug dosage information should not

be input to any workflow that contains clustering components unless it is generalized by at least 5-anonymity:

```
CONTEXT InLink(?l) ^ hasFile(?l,?d) ^ hasAttribute(?d,?a)
      ^ DrugDosageAttribute(?a)
PROTECTION
  REQ: -
  DIS: hasNode(?w,?n1) ^ hasComponent(?n1,?c)
      ^ Clustering(?c)
CORRECTIONS
  hasNode(?w,?n) ^ hasParamValue(?n,level,5)
      ^ hasOutLink(?n,?l) ^ hasComponent(?n,Anonymize)
```

Finally, we added a module to Wings to check compliance with a given set of privacy rules, after the workflow generation algorithm is completed.

Open Questions and Requirements

In this section we describe a number of open questions and requirements derived from our insights from these use cases.

A Usage-Oriented Policy Language

A language for representing privacy policies for workflows needs to be developed, together with a semantics for reasoning about it. The language needs to support a variety of aspects about private information and privacy relevant algorithms and support novel types of privacy policies, such as:

- Algorithmic policies, to specify what kinds of data analysis algorithms are allowed. These could be allowed for specified data types, for specific data sources, or in a data-independent manner. For example, group detection algorithms could be disallowed for use with medical data sources. Another example would be to disable the use of group detection followed by event detection algorithms unless the accuracy of the data sources is above a certain level. This policy may be used to avoid positive identification of individuals as threats with an accuracy so low that it may be a concern for individuals' liberties. Algorithmic policies may be contingent on properties of intermediate data products. Such policies may also express that certain steps have to be performed before storing a result, or transmitting data over an unsecured network. Expressing and reasoning about these types of policies may build on Linear Temporal Logic which has proved useful in other areas of computer science, most notably software verification and more recently automated planning (e.g., (Bacchus and Kabanza 1998)).
- Query-based policies, to specify what kinds of questions the system is allowed to act upon. These include both user-issued queries as well as system-generated intermediate sub-queries. For example, queries regarding payments may be allowed to the system in accessing any kind of sources including medical and financial sources, while any sub-queries regarding the nature or details of patient treatment may be disallowed.
- Data integration policies, to specify at the workflow level whether diverse data sources could be integrated through data mining steps. These would essentially control the legal joining of workflow strands.

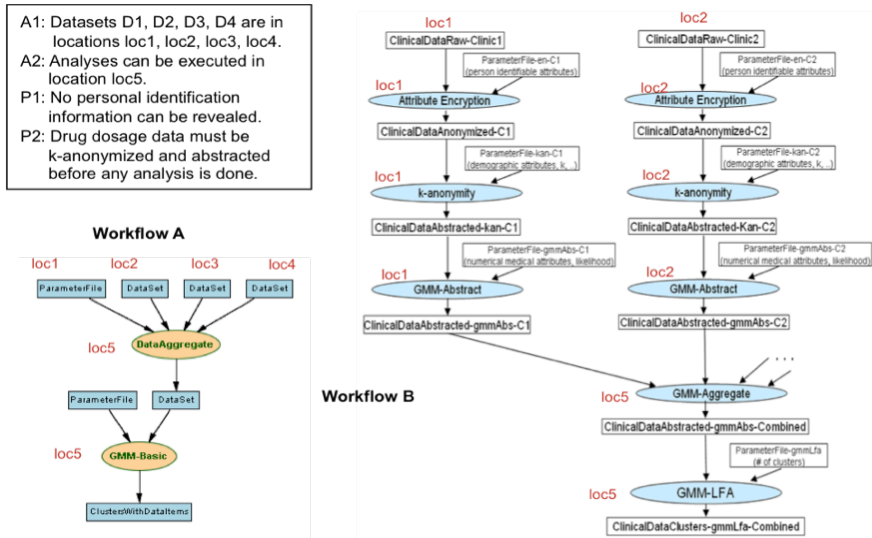


Figure 1: Two example workflows together with relevant assumptions (A1, A2) and privacy policies (P1, P2) (Gil et al. 2007a).

- Data creation policies, to specify what kinds of data may be created by the workflow. This could be specified via attribute types, entity types, or specific values.
- Provenance policies, to specify what information needs to be recorded and for how long it needs to be kept. This would reflect privacy needs for auditing and the stature of limitations for such requirements. Without these policies, there are no limits to the amount of details that a system could be expected to provide well after a workflow is used, so it is best to state these expectations up front.

These policies augment and are complementary to access policies for specific data sources or services in the system.

Extending Workflow Systems

Given this language, existing workflow systems would need to be extended in the following three ways.

1. Workflow creation and execution subsystem need to be extended. The workflow creation process that is responsible for selecting the data mining processes and data sources to be used in answering a query or line of inquiry needs to be governed by privacy policies that place constraints on the choices of data sources and algorithms. The extended workflow system should exercise full control over the design of the end-to-end data mining process before any computation occurs. The execution system needs to enforce privacy constraints that regard decisions about where data is being analyzed, and to enforce aspects that are only evaluable during execution itself. For example, a privacy policy may state that if the output of a clustering algorithm contains a cluster with less than k individuals then the analysis is not allowed. Generally the fidelity of the models of applied components will not be high enough to predict such situations ahead of execution.
2. Workflow systems need to leave detailed provenance trails of how data was processed and what mechanisms were used to ensure compliance with privacy policies by the workflow, both in its design and in its execution, in order

to support transparency and accountability regarding violation of privacy policies that regard the use of data. Re-execution of workflows through provenance trails could be used to prove, during an audit, that a given result was obtained as advertised.

3. Workflow system should support a distributed architecture for storage and retrieval of policy information. There may be several ways in which privacy requirements enter the system. Privacy rules need to be associated with different entities in the system. Some privacy policies should be associated with data when it is collected. Other privacy policies would be associated with collections or types of data (e.g., all the data collected by a clinical trial). Yet other policies may be application or system specific (e.g., federal or state privacy laws that may apply).

An important open issue is the trade-off between privacy and result quality. Many privacy preserving operations abstract information from the data which leads to less accurate results. Data descriptions and algorithm models will have to be extended to represent the relative accuracy of algorithms based on abstraction data features.

Reasoning about Privacy and Privacy Policies

An important open question is the negotiation of policies. Mechanisms need to be developed that support argumentation of “need to know” to relax privacy requirements if needed. When the privacy policies are too constraining for the system to find a solution to a query, it is possible to explore relaxations of some subset of policies that would enable the original request to be fulfilled. By articulating the choices that the system rejected and the privacy policies that forbid those analyses, the system would be articulating its “need to know” for specific data sources and data products. Conversely, the developed mechanisms could be used to check whether existing information disclosure agreements are indeed necessary for the purpose, or whether the level of privacy could be increased, e.g., via the inclusion of

additional anonymization steps, without adversely affecting the quality of the final result.

Such mechanisms for reasoning about policies may also assist in the design of privacy policies themselves, by enabling exploration of allowable but undesirable workflows under a given set of policies. This is important, because it may be difficult to design policies that are complete, in the sense that there is no way to exploit sensitive data when complying with them.

Conclusions

In this position paper we have argued for the need of technology to support reasoning about privacy policies that regard the *use* of data. We believe that computational workflow systems are a good starting point and could be extended to support a variety of privacy related tasks including:

Ensuring compliance of a data analysis system with specified privacy policies before enabling execution and during execution via monitoring.

Assisting users to comply with required privacy policies by selecting data analysis workflows that comply with those policies for the datasets to be analyzed.

Enabling transparency of data analysis systems that use sensitive information, including the generation of detailed provenance trails.

Supporting accountability with respect to the appropriate use of data in compliance with privacy policies.

Supporting negotiation and relaxation of privacy policies as well as access to data, by providing evidence for the “need to know” of sensitive data and, conversely, the ability to identify opportunities for an increase in privacy where such measures do not adversely affect quality.

Acknowledgments: We thank William K. Cheung, Varun Ratnakar, and Kai-kin Chan for their previous collaboration on this topic. Part of this research was supported by NSF grant: CCF-0725332.

References

- Agrawal, R., and Srikant, R. 2000. Privacy-preserving data mining. In *Proc. of the ACM SIGMOD Conference on Management of Data*, 439–450.
- Bacchus, F., and Kabanza, F. 1998. Planning for temporally extended goals. *Ann. Math. Artif. Intell.* 22(1-2):5–27.
- Berners-Lee, T.; Hall, W.; Hendler, J. A.; O’Hara, K.; Shadbolt, N.; and Weitzner, D. J. 2006. A Framework for Web Science. *Foundations and Trends in Web Science* 1(1):1–130.
- Cheung, W., and Gil, Y. 2007. Towards privacy aware data analysis workflows for e-science. In *Workshop on Semantic e-Science (SeS2007)*, at AAAI07, Vancouver, BC, Canada, July 22–26.
- Deelman, E.; Singh, G.; Su, M.; Blythe, J.; Gil, Y.; C., K.; Kim, J.; Mehta, G.; Vahi, K.; Berriman, G.; Good, J.; Laity, A.; Jacob, J.; and Katz, D. 2005. Pegasus: A framework for mapping complex scientific workflows onto distributed systems. *Scientific Programming* 3(13).
- Dugas, M.; Schoch, C.; Schnittger, S.; Kern, W.; Haferlach, T.; Messerer, D.; and Uberla, K. 2002. Impact of integrating clinical and genetic information. *Silico Biology* 2(34).
- Esponda, F.; Ackley, E.; Helman, P.; Jia, H.; and Forrest, S. 2006. Protecting data privacy through hard-to-reverse negative databases. In *Proc. Information Security Conference*, 72–84.
- Gil, Y.; Cheung, W. K.; Ratnakar, V.; and Chan, K. 2007a. Privacy enforcement in data analysis workflows. In *Proc. of the Workshop on Privacy Enforcement and Accountability with Semantics (PEAS’07)*.
- Gil, Y.; Ratnakar, V.; Deelman, E.; Mehta, G.; and Kim, J. 2007b. Wings for Pegasus: Creating large-scale scientific representations of computational workflows. In *Proc. of the Annual Conference on Innovative Applications of Artificial Intelligence*.
- Gil, Y.; Ratnakar, V.; Kim, J.; Gonzalez-Calero, P. A.; Groth, P.; Moody, J.; and Deelman, E. 2010. Wings: Intelligent workflow-based design of computational experiments. *IEEE Intelligent Systems*. To appear.
- Jiang, W., and Clifton., C. 2006. A secure distributed framework for achieving k-anonymity. *VLDB Journal* 15.
- Kargupta, H.; Datta, S.; Wang, Q.; and Sivakumar, K. 2005. Random data perturbation techniques and privacy-preserving data mining. *Knowledge and Information Systems* 7(4):387–414.
- Kohane, J., and Altman, R. 2005. Health-information altruists – a potentially critical resource. *New England Journal of Medicine* 353(19).
- Lindell, Y., and Pinkas, B. 2002. Privacy preserving data mining. *Journal of Cryptology* 15(3).
- Mandl, K.; Szolovits, P.; and Kohane, I. 2001. Public standards and patients’ control: How to keep electronic medical records accessible but private. *British Medical Journal* 322(7281):283–287.
- OECD. 1980. Guidelines on the protection of privacy and transborder flow of personal data. http://www.oecd.org/document/18/0,2340,en_2649_34255_1815186_1_1_1_1,00.html.
- P3P. 2002. The platform for privacy preferences 1.0 specification. W3C Recommendation.
- SAML. 2005. SAML, security assertion markup language v2.0. OASIS Standard.
- Sweeney, L. 2002. k-anonymity: A model for protecting privacy. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems* 10(5).
- Sweeney, L. 2004. Finding lists of people on the web. *ACM Computers and Society* 34(1).
- Taylor, I.; Deelman, E.; Gannon, D.; and M.S., S., eds. 2006. *Workflows for e-Science*. Springer Verlag.
- Uzuner, O.; Luo, Y.; and Szolovits, P. 2007. Evaluating the state-of-the-art in automatic de-identification. *J. of the American Medical Informatics Assoc.* 14(5).
- Weitzner, D.; Abelson, H.; Berners-Lee, T.; Hanson, C.; Hendler, J.; Kagal, L.; McGuinness, D.; Sussman, G.; and Waterman, K. 2006. Transparent accountable data mining: New strategies for privacy protection. Technical Report MIT-CSAIL-TR-2006-007, MIT.
- Weitzner, D.; Abelson, H.; Berners-Lee, T.; Feigenbaum, J.; Hendler, J.; and Sussman, G. J. 2008. Information accountability. *Communications of the ACM*.
- Weitzner, D. 2007. Beyond secrecy: New privacy protection strategies for open information spaces. *IEEE Internet Computing*.
- XACML. 2005. XACML, extensible access control markup language v2.0. OASIS Standard.