Aintno: Demonstration of Information Accountability on the Web

Joe Pato
Cloud and Security Research Lab
HP Labs & MIT CSAIL
Cambridge MA, USA
Email: joe.pato@hp.com

Sharon Paradesi, Ian Jacobi, Fuming Shih and Sam Wang
Decentralized Information Group
MIT Computer Science and Artificial Intelligence Lab
Cambridge MA, USA
Email: {paradesi, jacobi, fuming, samuelsw}@csail.mit.edu

***Abstract*—*Information Accountability* aims to encourage responsible use of information by combining clearly expressed usage policies with systems for detecting misuse, and offering the social tools to provide redress. Unlike conventional access control systems, accountable systems allow access to the data, but specify the purposes for which that access is allowed via policies. Information consumers, in such a system, are implicitly bound by a contract (social or legal) to those policies and should be able to provide an account of how the information owner's data were used. However, most of the current systems on the Web are not accountable.**

We have developed a simplified information accountability model and built an experimental platform that allows individuals to explore how their information may be at risk. We present a scenario in which a person denied insurance can explore and determine exactly why he was rejected and what actions he can take to prevent such an action in the future. A number of open problems remain, including identifying necessary incentives for participants in an information accountability setting; determining the consequences of applying excessive usage restrictions; coping with information originating for arbitrary sources; and applying context to understanding how to interpret information.

***Keywords*-accountability; privacy;**

I. Introduction

Social networking services are the siren song of the digital age – content is the currency of popularity, and to withhold information from your social circle is to fade into oblivion. Typically, however, information shared is information over which control is lost. Further, on the Web, sharing personal data is the equivalent of both living in a glass house and providing an arsenal of well-formed projectiles.

In an attempt to mitigate potential damage arising from information leakage, many systems implement access control models. Like locks on doors, the aim is to limit loss by controlling who has access.

This approach falls short in at least two ways: (1) it relies on our ability to anticipate who should have access; and (2) it assumes that an authorized user will make appropriate, compliant use of the information and won't use it to harm the information subject. Exclusive reliance on access control, however, is a fundamentally brittle approach in that it yields systems in which information, once revealed, is completely uncontrolled[6].

Thus begins a privacy conundrum. Sharing information can lead to greater social engagement, but by failing to adequately anticipate how an authorized recipient will use our information, we may suffer harm through the release or misuse of that information. Yet, choosing to remain on the sidelines and refraining from sharing information is also problematic. When the fear of exposure inhibits full participation in life, this, too, is a privacy problem.

Our group has been studying *information accountability* as a complement to access control for protecting information sharing. As described in Section 2, accountable systems provide the ability for decentralized systems to determine whether each use of data is/was permitted by the relevant rules for the particular person in the particular circumstance. They also make that decision available to access control, audit, and other technology for real-time enforcement or retrospective correction or redress. Information accountability enhances privacy protection through deterrence. The expectation of detection and redress inhibits data misuse and complements real-time access controls.

Shifting our focus to information accountability as the basis for considering information sharing and disclosure is more tractable than relying on abstract notions of privacy to address the conundrum above. Privacy is a difficult concept. Solove, for example, observes in *Understanding Privacy*[16], that *privacy* is a rather overloaded term and created a bottom-up taxonomy[1] to make addressing privacy issues more practical. However, privacy remains a complicated notion to address. Benjamin Wittes goes further and advocates changing the terms of the discussion. He suggests that *databuse* – the unjustified deployment of user data in a fashion adverse to the user's interests[20], is a more appropriate focus. Wittes observes that our concern with others manipulating our information is "not in proportion to whether that data is used in a fashion that protects our privacy or confidentiality but in proportion to whether it is used for our benefit or to our detriment and critically, how seriously to our detriment." Our approach to information accountability sets the focus on data use and misuse.

[1]Coarsely, Solove's taxonomy consists of four groups of activities: (1) information collection, (2) information processing, (3) information dissemination, and (4) invasion.

To ground the concepts associated with information accountability, we have built an experimentation platform. Using a motivating scenario (Section 3), we demonstrate how public information about individuals can be used to their detriment and how use of simple usage restrictions can limit this harm. The scenario is brought to life using our experimentation platform (Section 4). This platform allows the easy creation of new scenarios and allows individuals to examine how their own data on the Web may be used.

Our primary contributions are: (1) to shift the discussion of privacy and access control to a discussion of data misuse and accountability; (2) to identify open challenges for information accountability; and (3) to develop a tool that allows users to directly perceive how their data may be at risk.

II. Accountable Systems

Information accountability has two primary characteristics: information correctness and responsible use. Correctness is determined through examination of provenance associated with data – for example, how did it come into existence, how has it been transformed and why should I believe it. Responsibility refers to how information is used. Are information consumers adhering to the intentions of the information owners?

A. Components of an abstract system

Both aspects of information accountability are critical to our work, but for this paper we focus on responsible use – this corresponds to the Information Processing and Information Dissemination groups in Solove's privacy taxonomy. As described in a recent CACM paper[19], our concept of an accountable system – one that encourages responsible use of information – relies on three main components: the ability to express information use policies; the ability to monitor and reason over information use; and the ability to provide redress. Each of these areas presents technical and social challenges.

First is the ability to express data use intentions – both by data owners and data consumers. Data owners express how information should be used. Data consumers assert the purposes for which the information is being used. Technically, we face a challenge in choosing an appropriate language for expressing controls and intentions. Formal rule languages make correctness analyses and reasoning feasible, but are alien to how most people think. Compounding this problem is the social challenge of motivating concern for information protection. While many people express concern over potential privacy exposures when they share data, few read or understand privacy policies published on websites and misunderstand how technical mechanisms work, making it difficult to express their ultimate intent[13].

Universal transparency is the second component of an accountable system. A system is accountable when all transactions are automatically monitored, enabling simple – and automatic – detection of information misuse. This presumes that the purposes for which information is used can be compared against the policies defined by the information owner.

The third component is redress, providing the mechanisms to compensate information owners for harm arising from misuse of their data. Our aim is to inhibit misuse by raising the costs for those who violate the rules. Rather than imposing an a priori enforcement mechanism, we hold all actors, human and system agents, responsible for their behavior when they deviate from the rules that apply to the system.[2]

B. Simplifying the model

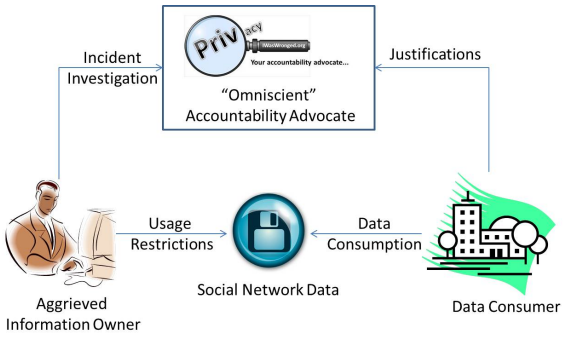

Figure 1. Simplified accountability model

As a first step towards achieving accountability, we have developed a simplified model shown in Figure 1. This model maintains the three components of our accountability architecture, but places limits on their functionality. First, rather than arbitrary policies for information use, we only allow simple usage restrictions to be placed on a user's data. By adopting an accessible approach similar to that used by the Creative Commons for describing information reuse rights[3], we remove the need for complex policy specification. These restrictions use the set of use annotations developed for the Respect My Privacy[9] project - identifying a set of purposes for which the information owners allow or disallow their information to be used.

Second, rather than assuming we can monitor all transactions, we rely on an *Accountability Advocacy* service, which arbitrates between the aggrieved user and the data consumer who may have misused the data. The Accountability Advocate need only have access to the data involved in a particular incident and the ability to interact with the data consumer. When a user believes that their data have been misused, they

[2]this is consistent with the view of information accountability outlined at the March 2008 NSF Cyber Trust Principal Investigators Meeting. http://www.cs.yale.edu/cybertrust08/Breakout-Accountability.pdf

[3]see: http://creativecommons.org

interact with the Accountability Advocate to investigate how their information was used.

Finally, rather than solely relying on law or regulations to compel data consumers to participate in our accountability environment, we assume that data consumers are willing participants in incident investigations, and we provide the Accountability Advocate, as a tool that simplifies their ability to do so. When an investigation is initiated by the information owner, the Accountability Advocate intercedes to request *justifications* from the data consumers explaining how information was used. If the data consumers are fully compliant with our accountability principles, they will honor all restrictions associated with the data. Even if they do not comply with these restrictions, we rely on them providing an accurate accounting of their information use as part of the justification provided to the advocate.

C. System Introduction

In our system, tools gather live data from social networks and the Web and are designed to allows individuals to explore how their own data may be used against them. We have created a data gathering engine to retrieve and derive facts from public data. These facts are available to information consumers and to the Accountability Advocate. We model the decision process at the information consumers using a formal rule language and a reasoner capable of producing formal justifications for a conclusion. This justification is processed into a form that can be explored by an interactive user using our user interface. The system is explained more fully in Section IV.

III. Scenario Walk-through

We have created a fictional character, *Danny Digger*, and allow interactive users to explore the consequences that befall him.

A. Motivating real-world examples

Our scenario is based on an increasing number of real-world examples where people's social networking information is used against them. Two of these incidents are briefly described here. In one incident, a Massachusetts Registry of Motor Vehicles employee was fired because he tweeted rants against his employer and customers on the microblogging website Twitter.[4] In the other incident, a woman was denied her teaching degree because she posted information about her students and co-workers to her MySpace webpage while on a student teaching placement. These postings violated the school's policy and among the items was a photo showing the student teacher wearing a pirate hat and drinking from a plastic cup that was considered to promote under-age drinking in "virtual view" of her students.[5] These cases highlight how personal information published on the Web can be used to harm an individual's interest.

B. Aintno Scenario

Danny Digger, our protagonist, is a chef at Hal's Wholesome Hotdog Hut who enjoys fine and casual dining. He shares information about himself through various social networks. Now he decides to apply for health insurance from the fictional insurance company *Aintno* but has been rejected.

Our demo simulates the use of public data by an insurance company making a policy application decision. Danny has been rejected for a new insurance policy due to poor dietary habits. Our system derives information on Danny's fondness for eating unhealthy foods from behavior observed in pictures and other postings on social networking sites.

In this scenario, the interactive user assumes the role of Danny Digger. The demo scenario opens with the user interacting with the Accountability Advocate. The Advocate, in our demo represented by the consumer self-help web site *IWasWronged.org*, allows the user to explore why Danny's insurance application has been denied. Aintno, the fictional insurance company, adheres to Information Accountability principles and cooperates with IWasWronged.org by allowing the user to examine a justification for Aintno's decision. The justification includes an explanation of the process and references to evidence that led to the decision on the insurance application. Some of this evidence may come from data obtained from Danny's Facebook or Flickr presence.

The interactive user is free to adjust data usage policies to make clear how he wishes social network data to be used. Once these policies have been expressed, he can re-run the simulation to see if other data will lead Aintno to the same conclusion or if excluding use of the data changes the insurance company's decision.

IV. Demo Architecture

The implementation architecture for the demonstration consists of a user interface, a data gathering engine, a reasoner, and a justification processing system. The implementation architecture for the demonstration platform is shown in Figure 2.

A. User Interface

Our demonstration platform is accessed via a web browser session starting at `http://dice.csail.mit.edu/aintno/ui`. The UI allows the end user to take on the role of Danny Digger and explore justifications for decisions made by consumers of Danny's data. A guided tour describing how to use the UI for the Aintno scenario is provided below in section V.

[4] http://www.myfoxboston.com/dpp/news/local/local-man-fired-for-twitter-rant-20110606 and http://www.bostonherald.com/news/regional/view.bg?articleid=1342870

[5] http://voices.washingtonpost.com/securityfix/Decision\%202008.12.03.pdf and http://abcnews.go.com/TheLaw/story?id=4791295

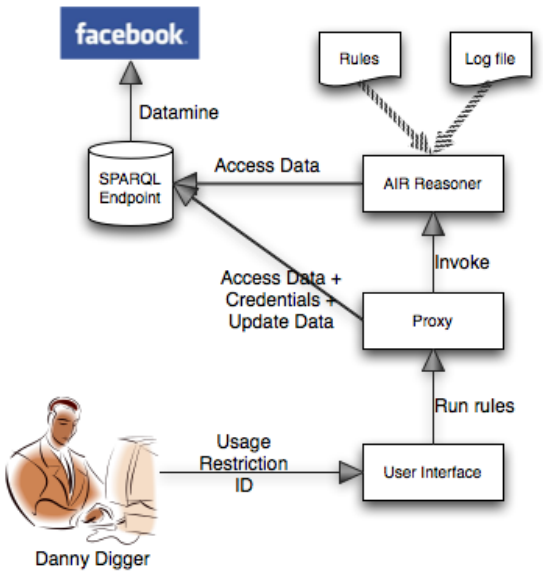

Figure 2. Implementation architecture

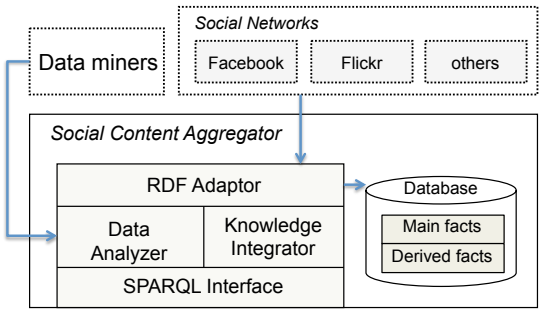

Figure 3. Data gathering engine

B. Data Preparation

Like many people, Danny has left behind digital fingerprints while using various web- and mobile-based applications. Major social network services like Facebook and Flickr host user data and find value by sharing data with each other[18]. These services often provide Web APIs for accessing the data programmatically. Therefore, data about Danny, although created in different places, are easily accessible and ready to be integrated to create a mosaic of his life. This linking and sharing of personal content from multiple sources makes it easy for data miners to make inferences about a user. Mining this content does not require a skillful programmer to break into the system and steal the data. GlobalInferencer[14], which also builds on our prototype data preparation system, illustrates this by making inferences about an individual's lifestyle and other behavior using linked data technology to perform unified searches across Facebook, Flickr, and public data sites.

As shown in Figure 3, our data-gathering engine consists of three parts: (1) a crawler that downloads data from Facebook and Flickr; (2) a social content aggregator that transforms data into RDF[12] format with vocabularies from predefined ontologies; and (3) data miners to assert newly inferred information from the subject's data. We use standard web APIs provided by Flickr and Facebook to allow our crawler to build a dossier about our fictional character Danny Digger.

First the crawler retrieves Danny's data such as photos, metadata of the photos, status updates, basic profile, and other personal content. To facilitate integration and interoperability of the fetched data, we created ontologies to describe the domain of social network platforms. The ontologies specify the relations between different types of information within the social network as well as possible references to data across platforms. This supports the example of Danny posting a link to his Flickr photo on his personal Facebook wall. Using the ontologies, the fetched data is transformed into RDF and saved to the back-end database with associated provenance information. Using Semantic Web technology for data representation, we can model complicated relations for personal data on the Web and support answering queries made in an RDF query language like SPARQL[3]. For example, we can respond to requests like "Show me all Danny's photos that are tagged with keywords including *hamburger*, *chicken wings* and *French fries*."

In addition to the data fetched from Facebook and Flickr, we use data miners to create *derived facts* that represent information inferred from the personal data collected about Danny. One of our naive data miners searched for the appearance of keywords like "hotdog" in Danny's photo tags to infer whether Danny has a strong preference for a certain kind of food. When the results show a high frequency of the keyword "hotdog", the data miner asserts a new derived fact `Danny eats hotdog` into the database. We include provenance information with each derived fact in the database to indicate from which source item it was inferred. For example, the derived fact `Danny eats hotdog` could be inferred from the Facebook album entitled *My favorite food*.

C. AIR Reasoner

In our model, a conforming data consumer is willing to provide the Accountability Advocate justifications for how a data subject's information has been used. Further, a fully compliant participant will honor any usage restrictions associated with the data. To automate production of these justifications and to control access to the data subject's information by the data consumer, we use AIR[10], an extension to N3Logic[2]. In previous work we have used AIR to support privacy protection in Web-based information systems[7].

For our scenario, we implemented a simple policy expressed as a set of rules which might be used to grant or deny someone health insurance coverage. The rules are designed to deny health insurance to anyone deemed to eat unhealthy food. Such a decision may be reached in one of several different ways as outlined by the following rules:

1) Eating a kind of food which is known to be unhealthy means one eats unhealthy food (i.e. instantiation of the class of unhealthy food products).
2) Eating at a restaurant which primarily serves a kind of food means that one eats that kind of food.
3) Working at a restaurant means that one eats at that restaurant.
4) Liking a restaurant means that a person eats the food served at that restaurant.
5) Liking a specific kind of food means that a person eats that kind of food.

Thus, according to the above rules, if it is known that:

- Danny works at Hal's Wholesome Hotdog Hut.
- Hal's Wholesome Hotdog Hut serves Hot Dogs.
- Hot Dogs are an unhealthy food.

then it should be possible to conclude that Danny eats unhealthy food.

An insurance company such as Aintno may be able to discover these facts in a number of ways, including through explicit information posted at social networking sites as well as statements derived from other media. For our system, we use the data gathering engine described above. When usage restrictions exist on some of these data, the reasoner relies only on facts not precluded for insurance purposes.

D. Justifications

The AIR reasoner generates justifications for any deductions made by the reasoner. These justifications include information about the structure of the deductions such as the ordering of nested rules, the facts triggering a rule, and the natural-language prompts and descriptions associated with a particular rule.

The justifications produced by the AIR reasoner, however, are not suitable for user presentation and must be translated for human consumption. The AIR reasoner produces an RDF graph which details every step of the reasoning that was performed over the relevant data and policy files. This RDF triple graph is not easy to visualize and present to users. In addition, since the AIR reasoner produces output expressed as URIs, there needs to be a component that translates these URIs into human-readable strings that can be presented to the end user.

The translation middleware converts the raw RDF graph into a data structure that is more suitable to present to users. The resulting translation is a directed acyclic graph (DAG) of nodes, where each node represents an application of a rule as defined in the policy. Parent nodes are logically dependent on their children. The root of the DAG is one node that either asserts that the given scenario is compliant with Aintno's policy, or is non-compliant with Aintno's policy.

A node can be dependent on another for multiple reasons:

1) A rule can directly cause another rule to fire. For example, a rule that states "If there is evidence that Danny likes hot dogs, then apply rule `SubjectLikesAFood`".
2) A rule can assert a fact that another rule requires to fire. For example, a rule that states "If there exists a photo of Danny eating a hot dog, assert fact `Danny eats hotdog`". A second rule which takes as input `Subject eats hotdog` is then logically dependent on the first rule.

These dependencies are spelled out in the RDF graph produced by the reasoner. However, the middleware translates the dependencies into explicit links between nodes in a DAG.

In addition, the middleware actively retrieves labels with which to print every triple that it encounters. We expect that data in this ecosystem will have associated `rdfs:label` triples which associate the URI with a human readable representation. The AIR reasoner does not care about labels associated with URIs, however, and does not record them in the justification. As part of the construction of the DAG, the middleware searches the included data log files, as well as the SPARQL store for appropriate labels for each of the referenced URIs.

The final output of this stage of the system is a DAG with natural-language text in each of the nodes describing each specific rule application present in the AIR justification. This DAG is suitable to be presented to an end-user in a web format.

V. Demo Guided Tour

When you first enter the demo, you arrive at a session on *IWasWronged.org* that is already in progress. Figure 4 depicts the initial page which is open to a tabbed folder for

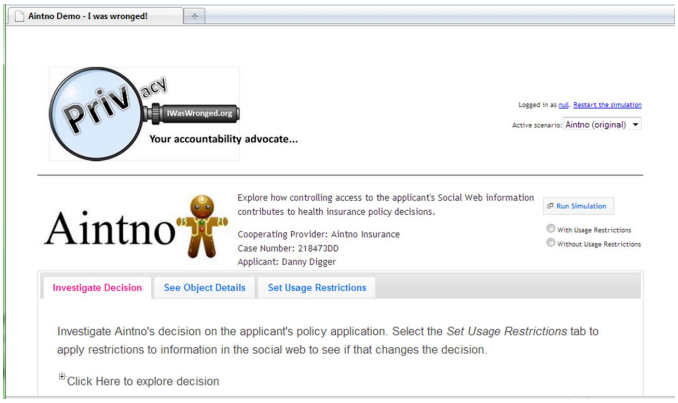

Figure 4. Overview of the demo – landing page

viewing the decision on Danny's insurance application. The *Investigate Decision* tab is open and you are prompted to click to explore the decision. The demo also allows you to see details of evidence used to support Aintno's decision in the *See Object Details* tab. Usage restrictions can be set on an object used in the decision in the Object Details tab or can be applied to collections of information in the *Set Usage Restrictions* tab.

1) Starting the Simulation: Exploring the decision provides a structured traversal of Aintno's justification for its decision. Clicking on the prompt text invokes the decision process – Aintno uses a set of rules describing conditions necessary for granting coverage as well as conditions that result in denial of the application. Aintno considers Danny's application as well as data it retrieves from Danny's public Facebook and Flickr pages. When using the initial demo conditions, Aintno determines that the policy application should be rejected.

The display uses a familiar file explorer display model. You can expand or contract elements of the explanation by clicking on a line. If an element of the justification is supported by more detail, a [+] icon appears on the left of the line and you can click this to disclose the supporting information. If you would like to hide detail, click on a line that is preceded by a [–] icon.

Figure 5 shows the decision explained through a series of questions and answers terminating at a set of facts that are evidence for a conclusion Aintno has made.

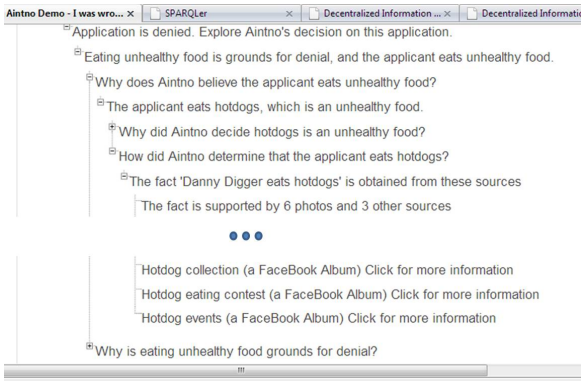

Figure 5. Justification for denial of Danny's insurance

2) Exploring evidence: Evidence used to support a conclusion is displayed as a simple statement in the hierarchical output. Typically, the information listed includes an active link to more detail about how that evidence was obtained. Clicking on an evidence line will open detail in another browsing tab or page. When the evidence is a picture, hovering over the picture will show some more detail which can be further explored.

3) Setting usage restrictions on evidence: Usage restrictions can be associated with information obtained from Facebook and Flickr. These restrictions use the set of annotations developed for the Respect My Privacy project, identifying a set of purposes enabled or precluded by the information owner. The demo supports setting the full range of purposes, but only the Insurance purpose will affect the decision Aintno makes regarding Danny's application.

Restrictions can be set in two ways – in the *See Object Details* tab or in the *Set Usage Restrictions tab*. They can be set for a given object when looking at the source detail in the *See Object Details* tab. This tab is activated when you request more information when hovering over an image displayed in a decision.

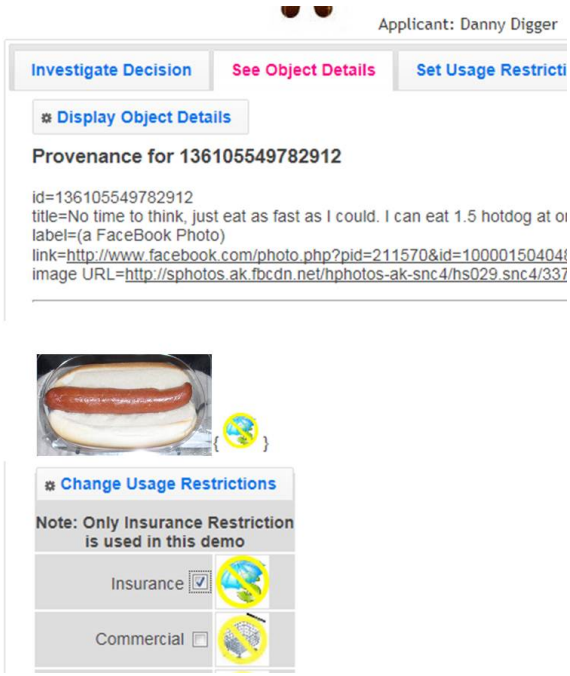

Figure 6. Provenance details

In Figure 6, detail for one of the images showing Danny eating hotdogs is displayed. The image is also shown as a thumbnail and the set of restrictions associated with the image are displayed within braces next to the thumbnail. A portion of the menu used for setting usage restrictions is also depicted above.

Usage restrictions can also be set and viewed from the *Set Usage Restrictions* tab. You can view several collections of images obtained from albums in the social network. Restrictions can be applied to all objects viewed in this tab – simply select the check box next to the image or album name and then scroll to the restriction setting menu below and choose which purposes should be restricted.

4) Honoring Restrictions: The demo allows the interactive user to choose Aintno's behavior with respect to Information Accountability principles – you can choose to have the company honor usage restrictions, or allow the

company to ignore any policy associated with the social network data. Unlike conventional access control systems, Information Accountability usage restrictions allow access to the data, but specify for which purposes that access is allowed. It is the responsibility of the data consumer to adhere to those policies and to be able to provide an accounting of how the data has been used.

To change Aintno's behavior regarding honoring restrictions when analyzing Danny's insurance application, select the *With Usage Restrictions* radio button in the upper right area of display screen and then click on the Run Simulation button immediately above the radio buttons. This will bring you to the *Investigate Decision* tab ready to rerun the simulation with Aintno honoring the usage restrictions you have set.

VI. Challenges for Information Accountability

Our accountability model, with its simplifications, have made it feasible to explore a prototype system, but several key problems remain.

A. Incentives

We have not explored realistic incentives for the key participants in our model. For data owners, the obvious path is to set restrictive usage policies on all of their information. When and why would a data owner relax usage restrictions? There must be a model where users benefit from having a data consumer access their information and as a result create tension over which restrictions to establish. Otherwise, the natural path is to place the most restrictive settings on all data.

For data consumers, it is unclear why they would adhere to our accountability principles. We must show how willing participation provides value for the data consumer – either through a direct economic incentive or indirectly, possibly by establishing good will and the creation of brand equity. If there is no benefit to voluntary compliance we may need to explore regulatory controls that establish penalties for deviation from the principles.

Finally, we need to establish the value of serving as an Accountability Advocate. Such an advocate could be operated by the private-sector if a business case can be constructed and the relying parties, the data subject and data consumer, trust the service operator. Alternatively, it may be preferable to explore the rationale for a not-for-profit service, akin to the Better Business Bureau, or outline the interest a governmental agency may have in fulfilling this role.

B. User Responsibility

How does expressing usage limitations relate to personal responsibility? Is an individual free to avoid the consequences of making irresponsible statements or for spreading misinformation by setting usage restrictions that preclude use of that information in any context of significance? In April 2011, the Twitter hashtag ***#NotIntendedToBeAFactualStatement*** emerged in response to false claims made during a floor debate by Senate Minority Whip Jon Kyl. Derived from the response provided by a spokesperson for Senator Kyl to CNN reporters, annotating a statement with the caveat *not intended to be a factual statement* has become grist for political satirists. Do usage restrictions provide the same escape valve as this caveat? Conversely, does the use of usage restrictions suggest that information shared by an individual should be suspect, treated as hyperbole or gossip because the individual is insulating themselves from harm associated with disclosure of the information?

C. Data Ownership and Provenance

Social networks are changing the way people perceive data ownership. On those networks, content that mention an individual is usually correlated with the actual user, either by tagging or commenting. Often these systems allow the data subject to control this correlation or refute a claim by removing the tag. On these networks there is an implicit acceptance of an individual's awareness of and agreement with publicly viewable personal information. In the broader Web, however, data about a person can originate from other individuals and is much harder to control. Three sources of external-entity data ownership are: (1) friends and family on social networking sites; (2) Companies that obtain anonymous user feedback (for example, product ratings on Amazon or movie ratings on Netflix); and, (3) public records such as house purchase records placed online by realtors or through a government registry of deeds.

In our system, we preserve relevant pieces of meta-data as the provenance for information we gather. This allows the information owner to more clearly understand the sources of facts used in a data consumer's decision. We would like to take this concern for information correctness and attribution further. As suggested by Aldeco-Perez[1], we would like to add details of data derivations involved in the inference process and use measures of data source quality to shape the conclusions our system generates.

D. Context

Context can be a critical tool for information accountability. As Solove observes,"[w]hen data is removed from the original context in which it was collected, it can more readily be misunderstood."[16] This leads us to a counter-intuitive reflection that sometimes it is better to share more information than less. For example, a cropped photo depicting one individual lunging at another can be fairly damning. The full image displaying a theater stage places the scene in context and changes the viewer's reaction. The challenge for information accountability is to develop techniques, possibly through a more complete provenance model as suggested above, to detect misuse through the use of information out of context.

Context is also important in determining when information should be shared. Madejski et al. [11] show the results of an evaluation that measured privacy attitudes and sharing intentions and compared them to actual privacy settings on Facebook. Their recommendations for Facebook align with our notion that privacy settings based on object-type are not sufficient, we need context-based privacy settings.

VII. Related Work

Access control or mediated access systems are traditionally used to limit access to sensitive information. A number of these systems have been applied to control access to the growing body of information individuals share about themselves on social networks. Kang et al. [8] describe the problems that arise when individuals subject themselves to self-surveillance – gathering information about themselves using mobile devices and sensors. They cite the difficulties for individuals to understand how to properly grant access to this data and recommend a Personal Data Guardian as the solution. The Personal Data Guardian is a service responsible for providing secure information storage and staffed by an employee who would maintain a Personal Data Vault for subscribers. This is similar to the system designed by Tootoonchian et al. [17], called Lockr, that separates personal information on social networking sources from other pieces of information. Such decoupling allows users control of their own personal data. It also acts as an intermediary between users and service providers, thereby providing an additional layer of security. Neither of these systems, however, deals directly with controls if information escapes the guardian, nor do they deal with information created and owned by a third party. For example, if Danny's local registrar of deeds places information about Danny's recent house purchase on the town's website, the owner of that content is the registrar of deeds and not Danny.

Sloan et al. [15] describe various technical and public policy challenges facing information accountability systems. They cite the role of informational norms as the unifying theme in ensuring adequate information privacy. They find the following key problems for developing effective accountability systems:

- Developing machine-readable forms of subtle, nuanced privacy rules.
- Ensuring the optimality of trade-offs made by privacy rules.
- Establishing new norms for information sharing
- Creating incentives for participation in accountability systems.
- Resolving inconsistencies in norms among different population groups.

Gajanayake et al. [5] propose solutions for sharing health care data using information accountability principles. Privacy and accountability are important in the health-care industry because of the sensitive nature of the information being passed around. As the authors point out, information accountability mechanisms have to identify the parties that can be held accountable, issues for which they can be held accountable and appropriate mechanisms for each party. One such mechanism could be detecting when an unidentified entity gains access to a patient's record and notifying the patient. The authors describe a system based on the Health Level 7 standards that incorporates three types of classes and associations for the entities under consideration.

VIII. Conclusion

Information Accountability aims to encourage responsible use of information by combining clearly expressed usage policies with systems for detecting misuse, and offering the social tools to provide redress. Shifting to accountability as the basis for considering information sharing and disclosure is more tractable than abstract notions of privacy and more appropriate than focusing solely on access control. By attacking the problem of information accountability, we get a glimpse of the technical and legal infrastructure that is needed to provide the social value people typically think of as privacy.

We have simplified our abstract accountability model to allow experimentation with some possible technical solutions. This has allowed us to create a tool that explores how public information can be used against an individual and can be used to increase public awareness regarding information misuse. However, a number of open problems remain, including identifying necessary incentives for participants in an information accountability context; determining the consequences of applying excessive usage restrictions; coping with information originating from arbitrary sources; and applying context to understanding how to interpret information.

Acknowledgment

The authors would like to thank Hal Abelson, Susan Landau, Mike Speciner, Jeri Zeder and members of the Decentralized Information Group for their input on this topic. This paper is based upon work supported by the National Science Foundation under Award No. CNS-0831442 and the Air Force Office of Scientic Research under Award No. FA9550-09-1-0152. We also thank HP and Qualcomm for their support.

References

[1] R. Aldeco-Perez and L. Moreau, "Information accountability supported by a provenance-based compliance framework," in UK e-Science All Hands Meeting, 2009.

[2] T. Berners-lee, D. Connolly, L. Kagal, Y. Scharf, and J. Hendler. "N3logic: A logical framework for the world wide web," Theory Pract. Log. Program., 8(3), 2008.

[3] K. G. Clark, L. Feigenbaum and E. Torres, "SPARQL Protocol for RDF," in http://www.w3.org/TR/rdf-sparql-protocol, 2008.

[4] D. Fensel, F. van Harmelen, B. Andersson, P. Brennan, H. Cunningham, E. Della Valle, F. Fischer, Z Huang, A. Kiryakov, T. Kyung-il Lee, L. Schooler, V. Tresp, F. Wesner, M. Witbrock and N. Zhong, "Towards LarKC: a platform for web-scale reasoning," Proceedings of the IEEE International Conference on Semantic Computing, 2008.

[5] R. Gajanayake, R. Iannella and T. Sahama, "Sharing with care: an information accountability perspective," IEEE Internet Computing, 2011.

[6] L. Kagal and H. Abelson, "Access control is an inadequate framework for privacy protection," W3C Privacy Workshop, 2010.

[7] L. Kagal, C. Hanson, and D. Weitzner. "Using dependency tracking to provide explanations for policy management," IEEE International Workshop on Policies for Distributed Systems and Networks, 2008.

[8] J. Kang, K. Shilton, J. Burke, D. Estrin and M. Hansen, "Self-surveillance privacy," 97 Iowa L. Rev., March 2012, in press.

[9] T. Kang, L. Kagal, "Enabling privacy-awareness in social networks," AAAI Spring Symposium Series, 2010.

[10] A. Khandelwal, J. Bao, L. Kagal, I. Jacobi, L. Ding, and J. Hendler. "Analyzing the AIR language: a Semantic Web (production) rule language," In Web Reasoning and Rule Systems, volume 6333 of Lecture Notes in Computer Science, pp 5872, 2010.

[11] M. Madejski, M. Johnson, and S. M. Bellovin,"The failure of online social network privacy settings," Technical Report CUCS-010-11, Department of Computer Science, Columbia University, 2011.

[12] F. Manola and E. Miller, "RDF primer," in http://www.w3.org/TR/2004/REC-rdf-primer-20040210, 2004.

[13] A. M. McDonald, "Footprints near the surf: individual privacy decisions in online contexts," 2010, in http://repository.cmu.edu/dissertations/7

[14] S. Paradesi and F. Shih, "GlobalInferencer: linking personal social content with data on the web," in ICWSM-11 Workshop on The Future of Social Web, 2011.

[15] R. H. Sloan and R. Warner, "Developing foundations for accountability systems: informational norms and context-sensitive judgments," in Proceedings of the 2010 Workshop on Governance of Technology, Information and Policies, 2010.

[16] D. Solove, Understanding Privacy, Harvard Univ Pr, 2009.

[17] A. Tootoonchian, S. Saroiu, Y. Ganjali and A. Wolman, "Lockr: better privacy for social networks," in Proceedings of the 5th international conference on Emerging networking experiments and technologies (CoNEXT '09), 2009, pp. 169-180

[18] A. Tsotsis, "Flickr dips its toes into social with Twitter and Facebook 'share this' features," in http://techcrunch.com/2011/03/30/ flickr-dips-its-toes-into-social-with-twitter-and-facebook-share-this-features, 2011.

[19] D. J. Weitzner, H. Abelson, T. Berners-Lee, J. Feigenbaum, J. Hendler, and G. J. Sussman, "Information accountability," in Communications of the ACM, June 2008, pp. 82-87.

[20] B. Wittes, Databuse: "Digital privacy and the mosaic," in Justice and Law, Legal Architecture for the War on Terror, Information Technology, 2011.