Sci Eng Ethics (2015) 21:941–966
DOI 10.1007/s11948-014-9564-6

ORIGINAL PAPER

Data Mining and Privacy of Social Network Sites' Users: Implications of the Data Mining Problem

Yeslam Al-Saggaf · Md Zahidul Islam

Received: 30 January 2014 / Accepted: 27 May 2014 / Published online: 12 June 2014

Abstract This paper explores the potential of data mining as a technique that could be used by malicious data miners to threaten the privacy of social network sites (SNS) users. It applies a data mining algorithm to a real dataset to provide empirically-based evidence of the ease with which characteristics about the SNS users can be discovered and used in a way that could invade their privacy. One major contribution of this article is the use of the decision forest data mining algorithm (SysFor) to the context of SNS, which does not only build a decision tree but rather a forest allowing the exploration of more logic rules from a dataset. One logic rule that SysFor built in this study, for example, revealed that anyone having a profile picture showing just the face or a picture showing a family is less likely to be lonely. Another contribution of this article is the discussion of the implications of the data mining problem for governments, businesses, developers and the SNS users themselves.

Keywords Data mining · Social network sites (SNS) · Privacy · Content analysis · Logic rules

Introduction

The literature is rich with studies about privacy in the context of social network sites (SNS). Researchers employ different data collection and data analysis techniques to

Y. Al-Saggaf (✉)
School of Computing and Mathematics, Charles Sturt University, Boorooma Street, Wagga Wagga, NSW 2678, Australia
e-mail: yalsaggaf@csu.edu.au

M. Z. Islam
Centre for Research in Complex Systems, School of Computing and Mathematics, Charles Sturt University, Bathurst, Australia

study SNS, including crawling and automated data extraction to collect data, and social network graph and sentiment analysis for data analysis (Catanese et al. 2011; Alim et al. 2011; Thelwall et al. 2010). There are also several studies in the literature that examine the threats to privacy from data mining (Oboler et al. 2012; Birrer 2005; Rubenstein et al. 2008). However, to the authors' knowledge, there is no study in the literature that uses a content analysis technique to collect data and generate a natural dataset and then apply a data mining algorithm to the dataset to demonstrate, using the data mining technique itself, how data mining can threaten SNS users' privacy. The aim of this article is to explore the potential of data mining as a technique that could be used by malicious data miners to threaten the privacy of SNS users. The aim will be achieved by (1) performing a content analysis technique to collect data and generate a natural dataset and (2) applying a data mining algorithm to this dataset to provide empirically-based evidence of the ease with which characteristics about the SNS users can be discovered and used in a way that could invade their privacy.

The article makes three contributions to the literature about SNS. First, the article discusses privacy in the literature from a philosophical perspective to offer a logical and rationalistic basis for the importance of privacy in people's lives. Philosophical accounts of privacy are not often empirically-based and technical accounts of privacy do not normally explain why privacy matters on moral grounds. This article capitalises on the strengths of both approaches. Second, following the application of a data mining algorithm on a natural dataset to show the ease at which characteristics about the SNS users can be discovered and used in a way that could invade their privacy, the article discusses the implications of the data mining problem for governments, businesses, developers and the SNS users themselves. The third, and arguably the most important, contribution is the use of the newly proposed decision forest data mining algorithm (SysFor) to the context of SNS, to demonstrate, using SysFor itself, how hidden relationships between attributes in a natural dataset can be discovered and used in a way that could invade users' privacy. To explore the logic rules and patterns that exist in a dataset, SysFor does not only build a decision tree but rather a forest (i.e. a set of decision trees), allowing the exploration of more logic rules from a dataset compared to the number of logic rules that can be explored by a single tree.

In this article, the study performed a content analysis of 616 Facebook user's profiles which resulted in a natural dataset that contained information on 45 attributes related to those profiles and then applied SysFor to show the ease at which characteristics about the SNS users can be discovered and used in a way that could invade their privacy. In this study SysFor built around 800 generated logic rules that were then carefully analysed (based on confidence and support) to identify those strong rules (i.e. those having high confidence and support) that may be used to compromise the privacy of a Facebook user. It is hoped this innovative approach will stimulate debate among methodologists interested in privacy of SNS.

There are at least three different ways that SNS users' information can become available to third parties, be it persons or organisations: through negligence of their privacy settings, through the use of application programming interface in SNS and via tracking technologies like HTTP cookies (Sar et al. 2012). But these third parties

are not only collecting massive amounts of data about SNS users; they are also mining them for the purpose of placing the users in new and non obvious classifications or categories that the users themselves might never have imagined existed (Tavani 2011). Tavani further argues that the problem is that the current privacy laws (worldwide) do not offer individuals any protection with regards to how information about these users obtained through data mining are subsequently used, particularly to make decisions about them in light of the categories these data mining activities discover.

With the number of active users on Facebook reaching one billion (Facebook 2012), out of two billion internet users worldwide, there is no doubt that social networking is one of the most popular activities online. On social network sites (SNS),[1] web users generate a personal profile typically through answering a series of questions relating to their gender, age, location, interests, school and work history, etc., much like composing a curriculum vitae. In addition to hosting textual personal information, users are encouraged to post a profile photo of themselves to their personal profile. Once a simple profile is generated from the answers to these personal questions and an image uploaded (usually identifying the individual user), users can then enhance and modify their SNS profile by uploading more images, videos, links to personal websites, and allowing popular third-party applications access to their profile.

While on one hand SNS encourage the disclosure of different kinds of sensitive information related to the user (Al-Saggaf 2011), on the other hand, the revelation of personal information on SNS profiles such as one's gender, age, contact information, location, profile photo, and relationship status, can carry with it the risk of this information being accessed and used by others in ways that can violate this user's privacy (Al-Saggaf and Islam 2012).

Before demonstrating how data mining can actually threaten the privacy of SNS users and discussing some implications of this problem, the article will first address the questions: What is privacy? What are the main theories of privacy? Why does privacy matter? What is data mining? How can it threaten privacy? What are the different data mining techniques that can be used by web data miners to acquire users' personal information from SNS? And how can these techniques threaten the privacy of SNS users.

Privacy: Definition and Importance

What is Privacy?

For Tavani, privacy is defined as the ability to restrict others access to and control over the flow of one's personal information, including the transfer and exchange of that information (Tavani 2011: 136–137). Traditionally, privacy has referred to the

[1] For the purpose of this paper, a social network site is a web-based service that allows individuals to "(1) construct a public or semi-public profile within a bounded system, (2) articulate a list of other users with whom they share a connection, and (3) view and traverse their list of connections and those made by others within the system" (Boyd and Ellison 2007: 211).

capacity of an individual or group to have control over information about themselves (Rachels 1975). As a result, privacy is often cast in terms of a 'right' entitled a person and is closely tied to normative values such as autonomy and dignity (van den Hoven 2008). While legal conceptions of privacy have often been restricted to interpretations of privacy in terms of its violation, as in 'nonintrusion', or in terms of individual liberty, as in 'freedom to act' (Moor 1990), legal and normative accounts of privacy have attempted to reflect the increasing impact of information technology on the ability of individuals and groups to share, access, and use all kinds of information about persons. However, some theorists argue that these traditional theories of privacy now inadequately reflect the changes information technology has made to certain environments and situations that in the past were not problematic for privacy (Nissenbaum 1997, 1998).

What are the Dominant Theories of Privacy?

Perhaps the most influential account of privacy originates from the work of the philosopher James Rachels (1975). Rachels argued that what makes privacy so important, even when it comes to everyday personal information and situations, is because it affords individuals the ability to selectively disclose information relating to themselves. This control and management over self-presentation is vital for individuals to be able to successfully create and maintain different kinds of personal relationships. In other words, our ability to navigate a variety of social interactions depends on our ability to control information about ourselves. Without privacy, the variety of relationships individuals can participate in would disintegrate. So basic a need is privacy that it is often simply assumed as an individual 'right'.

Philosophers and policy makers have continued to elaborate on Rachels' 'control of information' theory of privacy, particularly in response to the impact of information technology. Moor (1990, 1997) has argued that the concept of privacy is better framed in terms of protection from intrusion and information gathering rather than strictly as control of information, and has defended a 'restricted access' theory of privacy. Moor defends the notion that privacy is first and foremost about protection from information gathering by others. This differs from the traditional account of Rachels that privacy is fundamentally about having control over self-presentation in regards to personal information. Moor's account reflects the impact of information technology, an influence that was less relevant two decades earlier when Rachels formulated his account of privacy. Specifically, on Moor's account, privacy may be preserved when it comes to information that is computerised—Moor refers to this as "greased data", or information that "moves like lightening and is hard to hold on to"—by allowing "different people different levels of access for different kinds of information at different times" (Moor 1997: 31).

On the other hand, Nissenbaum (1998, 2004, 2010) proposed a different theory of privacy. According to this theory, 'the contextual integrity framework', there are two types of informational norms: (a) norms of appropriateness and (b) norms of distribution (Sar and Al-Saggaf 2014). Norms of appropriateness dictate the nature of information about an individual that is allowable to be revealed in a particular context. For example, it is ok for my general practitioner to inform my specialist of

a medical condition she suspects I have, but it is not ok for her to share this information with her husband. Norms of distribution govern the flow of information from one party to another i.e. whether or not that distribution of information respects contextual norms of information flow (Sar and Al-Saggaf 2014). Using the same example, it is ok for my general practitioner to email my specialist, my medical report, but it is not ok for her to email it to my employer. The contextual integrity of the flow of information is maintained when both kinds of norms are respected; otherwise, a breach of privacy occurs.

Along with recognising that we often expect privacy even in public, Nissenbaum's account of privacy as contextual integrity critiques the value of the traditional dichotomy of regarding some piece of information as being either 'private' or 'public' when conceptualising privacy or formulating privacy policy. Rather, what we should be focusing on is what constraints ought to be imposed on certain information flows within particular contexts. She argues that the primary worry about privacy violations through new modes of information transmission is the fact that privacy is about the context the information is in and the principles that guide the sender and the receiver of personal information (Nissenbaum 2004, 2010).

Does Privacy Matter?

One possible answer to the question 'What situations or environments should we consider normatively private?' is that when and where people have a right to privacy is something that is culturally determined (Moor 1990). But Nissenbaum's theory of contextual integrity gives a different answer: a right to privacy depends on context-relative informational norms (Nissenbaum 2010). Facebook founder, Mark Zuckerberg, commented during a 2010 address that recent changes to Facebook's privacy policy reflected the changing expectations of today's internet users regarding privacy and disclosure in online environments. One commentator interpreted Zuckerberg's comments as implying that the "age of privacy is over" (Kirkpatrick 2010). The Facebook founder's comments inferred that people are increasingly willing to share more and more information online, implying that privacy, at least in 'public' online environments, matters much less than it used to for many people, at least judging by their social networking practices, and that privacy indeed *should* no longer matter as much in an online context.

But does the ostensibly cavalier disclosure behaviour of online SNS users really reflect the fact that people no longer value autonomy or control over their personal information (Johnson 2010; Kirkpatrick 2010)? Privacy and control over personal information online is still very much an important personal and social norm according to some commentators (Johnson 2009, 2010; Kirkpatrick 2010). Danah Boyd, for example, argues that teenagers and adolescents, the group making up the majority SNS users, do care about privacy and even think of their SNS page as a space where they can and should be allowed to control their information (Johnson 2009). Indeed, a study by Al-Saggaf (2011) found exactly this result. While self-disclosure was also common among the female participants he studied, these participants appeared to be aware of the danger of displaying sensitive information on SNSs and as a result were very conscious about their privacy (Al-Saggaf 2011).

Data Mining and Threats to Privacy

Data mining, also known as Knowledge Discovery in Databases (KDD), is the process of searching through databases in order to discover patterns or relationships within large datasets (Tavani 1999, 2011). Data mining techniques help users in various activities such as classification, clustering, association rule mining, and summarisation of data, making it possible to describe trends within the data, extract meaningful information, and predict the value of future data (Islam and Giggins 2011; Islam 2012; Rahman and Islam 2011; Rahman et al. 2011).

Data mining is an established practice in many industries including marketing, advertising, finance, banking and insurance. The information generated from data mining techniques can provide people working within these industries with accurate 'profiles' of consumers and their purchasing behaviour allowing them to more effectively target customers. Data mining techniques have also been used successfully in many other areas such as effective water management (Khan et al. 2011), policing and counter-terrorism (Birrer 2005; Rubenstein et al. 2008).

Web data mining can be broken up into three main categories: web structure mining, web content mining, and web usage mining (Kosala and Blockeel 2000). Web structure mining is the technique of extracting the links and structure of websites in order to discover previously unknown relationships between web pages (Ting 2008). Web content mining, sometimes referred to as text mining, is the process of scanning and extracting the content of a web page, i.e. text, pictures, graphs, audio, video, and hyperlinks, to determine the relevance of the content to a search query (Kosala and Blockeel 2000). Web content mining is particularly useful in social network analysis due to the rich content environment of social network sites (Ting 2008). Web usage mining is a technique used to track and analyse the navigation behaviour of web users (Ting 2008). As the name suggests, web usage mining analyses how websites have been used by users and to determine what they are looking for online.

All personal information published on the web by users can be searched and mined. Using the web and publishing personal information online carries the undeniable risk that once something is published on the Web, it will always be available for a business intelligence tool or a data mining application somewhere in cyberspace (Laurent 2011). But data mining raises unique concerns about the informational privacy of web users. The information generated via data mining techniques effectively circumvents the normative protection of personal informational privacy. This is because the data gathered tends to be, firstly, publicly accessible, i.e. "nonconfidential" in nature, and, secondly, the new information generated consists of analyses of patterns and relationships that are merely 'inferred' or implied from the vast amounts of individual instances of personal information within larger datasets (Tavani 1999, 2011). As a result, the method of data mining does not seem to directly violate one's informational privacy.

However, issues of privacy, autonomy, and consent do arise in regards to what data miners go on to use the resultant information for. Since the information effectively becomes the intellectual property of the data miners themselves, that information can be profited from by being sold to third parties. Threat to

informational privacy from web data mining comes from this 'secondary use' of personal data (Tavani 2011).

'Profiling' is one major example in which personal data can be used in unexpected ways that potentially threaten an individual's privacy and control over their personal information (Hildebrandt 2009; Tavani 2011). Profiling is the result of data mining whereby new facts or associations about an individual are inferred from personal information that is gathered from the different contexts, usually via the web. Retrieving certain bits of information from the content of an individuals' personal homepage which, once extracted, are then used in a completely different context, often resulting in placing the individual in a category or group that might not necessarily apply to the individual and may be socially detrimental to them, is certainly problematic. A potential consequence is that data mining techniques can encourage individuals and businesses that make use of profiling to engage in discrimination against people (Hildebrandt 2009).

The profiling example shows how the secondary use of mined personal data involves a certain lack of transparency on behalf of the data miner and a lack of consent on behalf of web users, who lose a degree of autonomy over how their personal information, published on the web, is used. This translates into a threat to the privacy of individuals who are often unaware of the ways in which their personal data is collected and eventually employed. The relationship between data miners and web users is, therefore, a rather one-sided affair. As van Wel and Royakkers have expressed it, "while most of the benefits go to the web miners, the web users are facing the dangers of web-data mining" (van Wel and Royakkers 2004: 139).

Data Mining in SNS

A consequence of the incredible growth in web user activity within SNS, such as Facebook and Twitter, is that the gathering of information in SNS has become very profitable, a prime target for data miners (Krill 2011). It has also made SNS into platforms for more malicious applications and illegal activities.

Protection of personal information from data mining techniques on SNS have not kept pace with the rapid evolution of SNS environments, which now connect and share users information not only with search engines, but with a myriad of third-party applications and external websites, increasing the risk of personal information being made accessible through various methods of data gathering.

Data miners employ several techniques to extract data from SNS including the use of *Web crawling* (Bonneau et al. 2009; Catanese et al. 2011), which is the most common technique, *Phishing attacks* (Jagatic et al. 2007), *Third-party applications* (Felt and Evans 2008) and by *Creating false profiles* (Bonneau et al. 2009). *Web crawling* involves employing "crawling" scripts (or 'spiders') to gather and sort through Facebook public profile listings (Bonneau et al. 2009; Catanese et al. 2011). Facebook public profile listings, which carry basic personal information such as a user's name, photo, and network memberships, are constantly crawled by search engines that are encouraged by Facebook in order to facilitate searches of registered members (BBC News 2007). Crawling social network sites allows data miners to

quickly and easily build a database of relationships of users through the rapid indexing of Facebook listings, making it a very effective initial data mining technique given that, on Facebook, public listings are enabled by default and less than 1 % of users opt out of allowing their profile listing to be publicly accessible (Bonneau et al. 2009).

Social network analysis, or data mining of social networks, has become a flourishing commercial enterprise. "Big Data" (Oboler et al. 2012) is now a concept applied to the growing masses of unstructured personal information dumped on social networking sites, like Facebook and Twitter, and for the new ways in which data miners can gather intelligence from all this personal information for profit (Krill 2011). One example of profiting from information mined on SNS is advertising and marketing. Facebook makes in excess of a billion dollars in advertising revenue, putting great financial pressure on businesses to turn SNS users' personal information into something commercially valuable (Laurent 2011). A direct result has been the emergence of various types of personally targeted advertising throughout SNS, especially Facebook.

Facebook's "Beacon" feature was an example of a data collection application that deposits a cookie on the user's computer, which then tracks the user's web behaviour, (e.g. their online retail purchases), on the various Beacon partner sites (Nakashima 2007). The collected information is then sold to relevant companies to improve product marketing and sales. The downside was that the application violated Facebook users' personal privacy by often reporting on confidential online purchasing behaviour to other users within the individual's social network without their consent (Manjoo 2007; Nakashima 2007; Tavani 2011). A search of recent media coverage that reported on invasions of privacy through data mining in SNS reveals that, more often than not, the perpetrator of invasions of users' privacy through data mining in SNS, like Facebook, is Facebook itself. This happens usually through Facebook advertising and the data driving it (Harfoush 2011).

Given that SNS are used by hundreds of millions of people, it is safe to assume that they profoundly influence peoples' lives, most of the time without their knowledge of the implications. Some have suggested that the security of personal information and the preference for disclosure of most SNS users are incompatible (Edwards and Brown 2009). The victims of personal privacy violations brought about by data mining practices in SNS are usually those users who tend to publish large amounts of personal information while at the same time disregarding available privacy measures (Gross and Acquisti 2005; Young and Quan-Hasse 2009). Next, we demonstrate how data mining can threaten the privacy of SNS users.

The Empirical Study

Sample

As explained in Al-Saggaf and Nielsen (2014: 462-463), *prior to data collection, ethical approval for the current research was obtained from the University's Human Research Ethics Committee. Using the search engine provided by*

http://youropenbook.org *site, two coders entered one keyword each ('lonely' and 'connected'), which were matched to the status description of all publicly available Facebook profiles (youropenbook.org 2012). The site was a specialised search engine that allowed all publicly available Facebook profiles to be accessed depending on a search criterion. Any profile status matching any of these keywords was returned to the researchers. However, the coders analysed only the first 308 Facebook profiles returned from the search for each status (i.e. the first 308 profiles for those who indicated in their status they were feeling 'lonely' and the first 308 profiles for those who indicated in their status they were feeling 'connected'). This was done to ensure that an equal number of records for both of these groups were obtained*. Having an equal number of records from both groups/categories is helpful for data analysis especially using data mining techniques such as decision trees and decision forests, since this avoids the class imbalance classification problems where most of the records of a dataset fall in one category and only a few fall in another category/categories. While the findings of this study are also applicable to male Facebook users and any two opposing characteristics, such as 'happy' and 'sad', could have been used to generate the dataset, the reason the data collection and analysis focussed only on females and the notions of connectedness and loneliness is because this study is part of a larger project that aims to investigate the relationship between female loneliness and self-disclosure in SNS and the implication of the loss of women's privacy on their safety in cyberspace. *In total, data were collected from 616 Facebook accounts which were retrievable via youropenbook.org (youropenbook.org 2012) at the time of data collection.* All records were for users who indicated in their Facebook profiles that they were females over the age of 18.

Content Analysis and the Creation of the Dataset

As explained in Al-Saggaf and Nielsen (2014: 463), *after a blank Facebook profile was studied to decide on the pieces of information to be included in the content analysis (Nosko et al. 2010), a total of 45 items were inserted in an Excel spreadsheet (in the columns) representing the pieces of information to be gathered from the selected public profiles and recorded in the spreadsheet. Two coders then independently performed content analysis by opening the selected public profiles from within their Facebook accounts and recording the information and the outcome of their analysis in the spreadsheet. Other than information about age, 'about', gender, city, relationship status, number of friends, likes, photos and interests, in all other fields only the presence or absence of information was recorded.* That is, only P (present) or A (absent) under the information cited was recorded. While the first coder analysed the 'lonely' profiles, the second coder analysed the 'connected' profiles. *The 308 female Facebook users were categorised as 'connected' based on a clear indication of this feeling in their latest wall posting at the time of data collection, such as one female who said "just had to share this message from Amber. We have never met, yet in a way I feel very connected to you". Conversely, the remaining 308 users were categorised as 'lonely' based also on clearly indicating this feeling in their latest wall posting at the time of data*

collection. For example, one female said "Such a boring evening…m so lonely…:-(". To limit subjectivity in the analysis, reliability was conducted on 62 profiles (31 profiles from the 'lonely' group and a further 31 profiles from the 'connected' group) representing 10 % of the sample. Percent Agreement of 95.2 % (N Agreements = 59) was recorded and a Cohen's Kappa score of 0.903 was returned indicating high inter-coder reliability. Disagreement between coders was resolved through discussion until a consensus was reached.

From the collected data we created a dataset which is a two dimensional table containing rows and columns. A row/record represents a Facebook user and a column/attribute represents the information of the users. For each user (i.e. for each row/record) the dataset contains various information (i.e. columns/attributes), out of which one is called the "class attribute" or "label" for a user. The value of the class attribute for our dataset can be either "connected" or "lonely" as explained before. That is, there are altogether 308 records with class attribute value equal to "connected", and other 308 records with class attribute value equal to "lonely".

The Application of the Data Mining Technique

By using data mining techniques a classifier (such as a decision tree, decision forest, artificial neural network and support vector machine) can be built which can study a dataset and learn the patterns for the records in order to be able to classify a record into one of the 'class values'. The classifier can then predict the 'class value' of an 'unlabelled' record for which the label is unknown. In this study, we built a decision forest by applying a recently proposed algorithm called SysFor (Islam and Giggins 2011) on our newly created dataset. A decision forest is a set of decision trees, that is, instead of building one decision tree SysFor builds a number of decision trees. In this study, we built 20 trees using SysFor. A decision tree discovers some patterns (i.e. not all patterns) of a dataset, whereas a decision forest can discover many more patterns, since it is a collection of a number of different trees that are obtained from the dataset. Therefore, a decision forest is generally considered as a powerful classifier. In this study, the class attribute in our dataset has two values (connected and lonely). Therefore, the decision forest that we built can classify a future 'unlabelled' record into either 'connected' or 'lonely'.

In order to explain the data mining techniques used in this study, we now briefly discuss two algorithms called C4.5 (Quinlan 1993; Islam 2012) and SysFor (Islam and Giggins 2011). While C4.5 was not directly used in this study, given that SysFor uses C4.5, an explanation of SysFor alone is not sufficient. For this reason we first explain C4.5. The C4.5 algorithm builds a decision tree as shown in Fig. 1. It is applied on a dataset having a number of records, non-class attributes (numerical and/or categorical) and a class attribute, where a record can have one of the possible class values for the class attribute. A record of our dataset has the class value either "connected" or "lonely". Numerical values (such as 1, 2 and 3) have natural ordering, while categorical values (such as "Sydney", "Melbourne" and "Brisbane") do not have natural ordering in them.

C4.5 aims to find a combination of attribute values so that, ideally, all records having the combination have the same class value. It first identifies the non-class

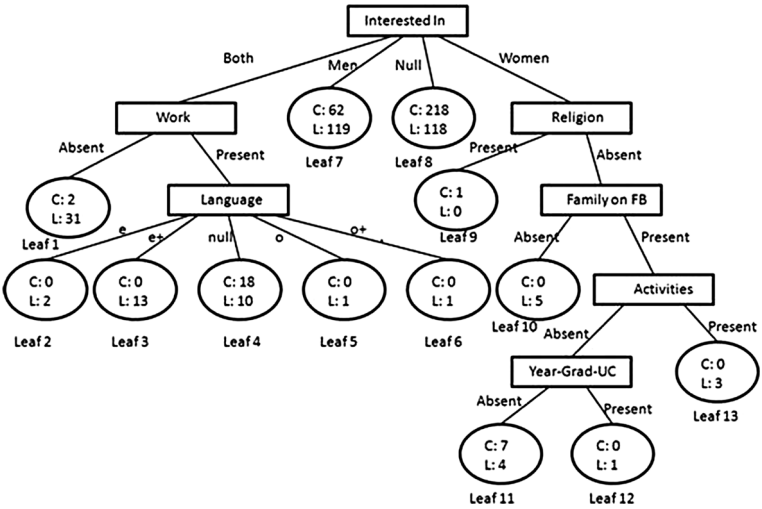

Fig. 1 Decision Tree-1 built from our SNS dataset

attribute that has the best ability to divide the dataset into mutually exclusive horizontal segments, where in each segment all records ideally have the same class value. In order to identify the best non-class attribute, it divides a dataset into horizontal segments based on the values of a non-class attribute so that each segment has the same value for the non-class attribute, if the non-class attribute is categorical. Therefore, the number of segments is the same as the domain size of a categorical non-class attribute. However, for a numerical non-class attribute the dataset is divided into two mutually exclusive horizontal segments, where all records in one segment have values greater than a constant 'c' for the numerical non-class attribute, and the records in the other segment have values less than or equal to c. The constant c is determined in a way so that it gives the best result for the attribute. Now, the records within each segment may have different class values. The homogeneity of the distribution of the class values within each segment is then assessed using measures called "entropy", "gain" and "gain ratio".

This process is repeated for all non-class attributes. The non-class attribute having the most homogeneous distribution within each segment is considered to be the best attribute and is chosen to be the root attribute of a tree. For example, in Fig. 1 the attribute "Interested In" is chosen as the root attribute.

Once a suitable root attribute is chosen the dataset is divided into horizontal segments based on the values of the root attribute. If the root attribute is categorical then the dataset is divided into as many segments as the domain size of the attribute, but if the root attribute is numerical then the dataset is divided into two segments using the constant c. The root attribute in Fig. 1 is a categorical attribute and therefore it divides the dataset into four segments where in the first segment all records have the value "Both" for the attribute.

C4.5 then repeats the whole process in each segment and finds the best attribute (same as how it finds the root attribute) within the segment. For example, in Fig. 1

C4.5 finds the attribute "Work" as the best attribute within the segment where all records have the value "Both" for the root attribute "Interested In". The process continues until it reaches a segment where either it fails to find any attribute that qualifies to be the best attribute within the segment or all records of the segment have the same class value. At the end of the process we get a complete decision tree.

SysFor (Islam and Giggins 2011) builds a decision forest, i.e. a set of decision trees, instead of one single decision tree from a dataset, based on a user input on the number of trees. For building a tree it uses an existing decision tree algorithm such as C4.5 (Quinlan 1993) and Explore (Islam 2012). If it uses C4.5, it builds the first tree as explained above. It then builds the second tree by using the second best "candidate attribute" for root (the best "candidate attribute" is used as the root while building the 1st tree) and then builds the rest of the second tree by using the same approach as explained above for the first tree, except that the root attribute is now different and therefore the subsequent data segments are also different. Similarly it builds the third tree and so on until either it uses all the high quality candidate attributes for the root or it completes building the user defined number of trees. Note that some attributes of a dataset may not satisfy the minimum quality requirement (assessed by measures called "goodness threshold" and "gain ratio") for being considered as a high quality candidate attribute. If it runs out of high quality candidate attributes for the root node and still requires building more trees (based on the user defined number of trees) then it uses the high quality candidate attributes at the second level of the best trees. For example, it can use the second best attribute (instead of the attribute "Work") at Level 2 where it uses "Work" in the 1st tree (see Fig. 1).

An obvious advantage of a forest is that it can explore more logic rules from a dataset compared to the number of logic rules that can be explored by a single tree. Since a dataset typically contains more valid patterns (i.e. logic rules) than the number of patterns that a single tree can discover, it is often worthwhile to build a forest in order to get a better understanding of the patterns existing in a dataset.

Once a decision tree is built then it can be used to predict the class value of a new unlabelled record for which the class value is unknown. For example, if we get a new unlabelled record having the value "Both" for the attribute "Interested In" and "Absent" for "Work" then we can use the tree shown in Fig. 1 to predict the class value of the new record to be "Lonely" since (as discovered by the tree) there are 33 records (31 + 2) in our dataset having the combination of the attribute values and 31 of them have "Lonely" as the class value. It is also indicated in the literature that typically a decision forest can make a more accurate prediction than a decision tree (Islam and Giggins 2011).

Results

We initially collected information on 45 attributes for 616 Facebook accounts. We then pre-processed/prepared a dataset with 22 non-class attributes and 1 class attribute with the aim of applying SysFor to explore the relationship between the non-class attributes and the class attribute, i.e. to explore the patterns for the people

feeling lonely and connected. We processed the collected data by dropping some attributes (from the original dataset having 45 attributes) which may not have any/significant influence on the class values. For example, the attribute "Gender" has only one value (i.e. "female") for all records since we only collected data for the female Facebook users. The attribute is not useful for the classification task since both "lonely" and "connected" users have the same value ("female") for the attribute. The reason why the information on the attribute was collected in the first place is to use the dataset also for a bigger study.

SysFor algorithm expects an attribute to be either numerical (all values for the attribute to be numerical over the whole dataset) or categorical, but not a mixture of numerical and categorical values for the same attribute. Therefore, during the pre-processing, some attribute values were also carefully categorised to suit the algorithm. For example the attribute "Number of Albums" had numerical values (such as 1, 2, and 3 based on the actual number of albums of a user), but for some users/records the attribute value was not accessible (due to their privacy set up), which is not the same thing as having the value equal to zero. We therefore converted the attribute into four sensible categorical values: None (if the value is inaccessible), Low (if the value is less than 5), Medium (if the value is between 6 and 20) and High (if the value is greater than 20).

The non-class attributes used in the pre-processed dataset are Profile Image, Relationship Status, Interested In, Family on FB, Hometown, Sex, High School, Year Graduated from High School, University/College, Year Graduated from Univ, Timeline, Works at, Friends, Number of Albums, Number of Photos, Language, Religion, Activities, Email, DOB, Political views, and People Who Inspire You. Of course the class attribute is "Status", which has two values; "lonely" and "connected".

The domain values of the attributes are also chosen carefully to be meaningful to the researchers in the context of the study. For example, the values of the attribute "Profile Image" are considered to be "Image shows user alone", "Shows user with one or more friends", "Shows the user at a special occasion", "Shows the user with their romantic partner", "Shows the user smiling", "Shows the user in a unique location that is not their hometown", "Face/head shot of the user", "Shows the user playing or watching sport", "Shows the user with their family", "Depicts an object with apparent meaning to the user", "Described as having a unique visual effect", and "Photo with a lot of exposure". Similarly, values for the attribute DOB are "Shows Full DOB", "Shows just Day and Month" and "Absent". However, many attributes such as Family on FB, Works At, and Hometown have only "p" (for present) and "a" (for absent) as their domain values.

We ran the SysFor algorithm on the dataset and built a decision forest with 20 decision trees. Figure 1 shows the 1st (out of 20) decision trees, where the rectangles are nodes and the ovals/circles are leaves. Each leaf can be represented by a logic rule. For example, Leaf 1 of the 1st tree can be represented by the logic rule "Interested In = Both and Work = Absent → Lonely (33/2)",[2] which suggests that in the dataset there are altogether 33 users who are interested in

[2] This is how logic rules are textually represented in data mining literature.

both male and female (as represented by Interested In = Both), and do not provide any work related information (as represented by Work = Absent). Out of these 33 users 31 are "lonely" while the remaining 2 are "connected" as represented by "Lonely (33/2)". Tree 1 (see Fig. 1) has 13 leaves and therefore 13 logic rules.

We also present five out of 20 trees in a textual form in the Appendix below. The trees are produced by our programs, implementing the SysFor algorithm, in the textual form. The presentation of five trees in the original textual form may give a better idea about the trees that are generated. Note that these are not the first five trees generated by SysFor, instead we present five trees with reasonable size for better understanding of the logic rules. Tree 1 has also been presented in a graphical form in Fig. 1. By "Level (0) Interested In" in the textual presentation of Tree 1 it means that the attribute "Interested In" is tested at Level 0, i.e. at the root level as we can see from Fig. 1. By "Level (1) Work when b" it means that when Interested In = b (i.e. both) then the attribute tested at Level 1 is Work (see Fig. 1). By "Level (2) {c;2, l;31} when a" it means that when Work = absent, then we reach a leaf having 33 records, out of which 31 are lonely (l) and 2 are connected (c). Therefore, the logic rule for the 1st leaf is "Interested In = Both and Work = Absent → Lonely (33/2)". Similarly, the logic rule for the 2nd leaf is "Interested In = Both and Work = Present and Language = English → Lonely (2/0)". Here, p and e represents Present and English respectively.

In the whole decision forest (i.e. 20 trees), we have more than 800 logic rules that describe various patterns for a user to be identified/classified as either "lonely" or "connected". Out of more than 800 discovered logic rules, we next discuss a few interesting rules expressed by RID (i.e. Rule ID). We mark a logic rule as interesting based on its 'support' and 'confidence'; two commonly used terms in data mining that are used to validate a logic rule. For example, the 'support' of the logic rule "Interested In = Both and Work = Absent → Lonely (33/2)" is $\frac{33}{616}$, since there are 33 records (out of total 616 records) where a user is interested in "both" and do not provide work related information (i.e. Interested In = Both and Work = Absent). Additionally, the confidence of the rule is $\frac{(33-2)}{33}$, since out of the 33 records 31 of them are "Lonely", which is the predicted class of the rule. High support and high confidence make a logic rule strong. A strong logic rule revealing a non-obvious and new pattern can be 'informally' considered to be interesting.

Note that, we also have many interesting logic rules other than those presented below. We present some interesting logic rules to demonstrate that it is possible to obtain some logic rules (by building a forest on a Facebook dataset) that can then be used by a malicious data miner to threaten the privacy of a user. If this can be explained by the logic rules below, then the existence of more sensitive-rules can only strengthen the argument. Rule ID 1 (RID 1) is taken from Tree 1 (see Leaf 1 of Fig. 1), whereas all other rules below are taken from other trees.

RID 1: Interested In = "Both" and Work = "Absent" → Lonely (33/2)
RID 2: Interested In = "Both" and Language ≠ "Absent" → Lonely (48/2)
RID 3: Interested In = "Both" and Language = "English Plus Others" and Home Town = "Present" → Lonely (24/0)

RID 4: Language = "English plus Others" and Activities = "Present" and Political View = "Absent" → Lonely (41/3)
RID 5: Interested In = "Both" and Email Address = "Absent" and Work = "Absent" → Lonely (30/2)
RID 6: Relationship Status = "In a relationship" and Language = "English Plus Others" → Lonely (12/0)
RID 7: Profile Image = "Either shows the face/head of the user OR a family photo" → Connected (46/0)
RID 8: Profile Image = "A romantic photo of the user with a partner" and Political View = "Absent", and Religion = "Absent" and Work = "Absent" → Connected (21/0)
RID 9: Profile Image = "A romantic photo of the user with a partner" and Friends Number = "Between 100 and 500" → Connected (21/2)

The logic rules reveal interesting patterns obtained from the collected dataset. For example, RID 3 reveals that all 24 users who are interested in both male and female, speak English and other language/s, and provide information on their hometowns feel lonely. Similarly, RID 6 shows that all 12 users mentioning "in a relationship" in their relationship status, and who speak English and other language/s feel lonely. RID 5 states that 28 out of 30 users who are interested in both, and who do not provide us with their email address and work information, feel lonely.

Rule ID (RID) 7 reveals that anyone having a profile picture showing just the face or a picture showing a family is likely to feel connected. There are 46 of such users (having a face or family photo as the profile picture) in our dataset and all of them feel "connected". Similarly, according to RID 8, all 21 users having a romantic profile picture with their partners, providing no information on political views, religion and work, feel connected. RID 9 shows that 19 out of 21 users having a romantic profile picture with their partners and the number of friends between 100 and 500 feel connected.

RID 2 is another interesting rule that discovers that 46 out of 48 users that have an interest in both male and female, and do not disclose any information related to their language, feel lonely. Similarly, according to RID 4, users who speak English and other languages and present no information on political views and activities feel lonely. There are 41 such users out of which 38 of them feel lonely. Note that here "users" means female users, since our dataset was created by taking the profile information of female users only.

Discussion

A malicious data miner can follow the same process as that used in this study to build a dataset and learn a set of interesting logic rules. He/she can then apply the rules on female users who are not in his/her dataset and who have not posted any information on either being "lonely" or "connected". The data miner first needs to collect all necessary profile information of a user and then apply the logic rules in order to classify her either as "lonely" or "connected". For example, following RID

7 the data miner will classify any user having a family photo as the profile picture as "connected". Similarly, any user having an interest in both male and female and not disclosing any information on language in her profile can be classified as "lonely" with high probability, according to RID 2. Learning this additional information ("lonely" or "connected") of a user can be considered a privacy threat given that the user did not have control over their information. This also shows that SNS failed to protect the gathering of information by others, since the user originally did not disclose this information. Moreover, based on this additional information, Facebook users can also be approached by the data miner with different purposes, including match making and direct marketing, resulting in inconveniences and breaches of individual privacy mainly caused by the use of Facebook.

The above is only one example of how a data miner can create a dataset from Facebook user profiles, build a set of logic rules and then use the rules to classify other users who have not even revealed information on the class values, which are in this case "lonely" and "connected". The data miner can also build logic rules for any set of class values, other than lonely and connected. That is, a data miner can build a decision forest considering any other attribute (for example, "Interested In") as the class attribute and thereby learn the logic rules to classify a new user as "interested in male", "interested in female" or "interested in both". This ability to build a decision forest considering any other attribute as the class attribute and thereby predicting the class value of a user can create problems for users.

There are other examples of how a malicious data miner can get access to sensitive information about an SNS user through the use of data mining. A data miner can visit the SNS profiles of all the friends of an SNS user. Based on the information of each friend's profile, the data miner can build a record of a dataset where all records represent the friends of the SNS user. The data miner can then apply *a clustering algorithm* on the dataset in order to find groups of similar friends to the SNS user. If the data miner finds that a large cluster of records displays a common connection with a particular type of music, movies and religion, then even if the SNS user never indicated any interest in these things, the data miner can assume that the user is also connected with them (i.e. 'music, movies and religion) simply because a large group of his/her friends do have connections with them. This may disturb the user in many ways including causing problems for the user when applying for a job or worse subjecting him/her to extortion or stalking; but this of course will depend on the motives of the data miner (e.g. one with a criminal intent versus one mining data to more effectively target customers).

The above raises an important question: what stops malicious data miners from mining the data they acquire from SNS and using the results of data mining activities to identify users or link them with groups they don't necessarily belong to and in ways that can disadvantage or harm these users? Data protection to prevent an illegitimate use of the data is a line of reasoning that appears to be logical. It is true that, on the one hand, users seem to be aware of the fact that the personal information they disclose in SNS can be used in ways that breach their privacy or used against their interests (Boyd and Ellison 2007; Al-Saggaf 2011). On the other hand, they appear incapable of refraining from disclosing their sensitive personal information in the SNS they join (Edwards and Brown 2009). However, just because

users cannot help but go to SNS, it does not mean that SNS should not allow people to do whatever they want with their information, although SNS owners could say join our sites at your own risk.

The challenge facing data protection is that individuals, for example in Australia, are not implicated by privacy laws.[3] In other words, the privacy laws don't implicate the public in their individual capacity. That is, the laws do not protect an individual's privacy if another individual breaches it. Individuals can be taken to court and tried in relation to privacy offenses, but only under, for example, breaches of confidence, or defamation laws, i.e. not under privacy laws. In addition, Australian privacy laws subject only businesses that have a turnover of at least three million dollars, which means even 95 % of Australian businesses are not bound by these laws.[4] We discuss privacy laws and data mining in more detail in the next section.

Implications of the Data Mining Problem

Implications for Governments and Businesses

Privacy Laws and Data Mining in SNS

The protection of personal information of SNS users from data mining is not explicitly stated within most privacy laws (Tavani 2011). Most privacy laws of government agencies and other institutions discuss the collection and use of personal data in general without mentioning data mining in SNS (Al-Saggaf 2012), reflecting the public and private sectors' lack of awareness of the threats of data mining in SNS or the unimportance of this issue. In addition, mandatory government rules covering a wide range of activities relating to the collection and use of personal data by government agencies, institutions and companies, tend to overlap when it comes to their basic principles (Al-Saggaf 2012).

The NSW Privacy and Personal Information Act (PPIP Act), established in 1998, includes twelve Information Protection Principles (IPPs) that describe rules to which government agencies must adhere when collecting, storing, accessing, using or disclosing personal information. The key principles can be reduced to the following, (1) information must be collected for a lawful purpose that is also directly related to the agencies activities, (2) agencies must disclose why the information is being collected and what it is being used for, (3) agencies can only use the personal information for the purpose for which it was collected, for a directly related purpose, or for a purpose to which the individual has given consent, and (4) agencies can only disclose personal information with the individual's consent (PPIP Act 1998).

The information protection principles outlined in the PPIP Act are almost identical to those described in the UK's Data Protection Act 1998. The key

[3] Australian Government office of the Australian Information professional. 2012. Business and Me. Retrieved March 9, 2012 from http://www.privacy.gov.au/individuals/business.

[4] Personal communication with The Victorian privacy commissioner, Ms Helen Versey, on the 13 of February 2012 during the sixth Australian Institute of Computer ethics conference in Melbourne. The Victorian privacy commissioner gave the key note address at this conference.

principles are also strikingly similar to the government rules that regulate the data collection activities of European companies, as Caudill and Murphy (2000: 12) reveal: First, a company should have a legitimate and clearly defined purpose to collect information. Second, that purpose must be disclosed to the person from whom the company is collecting information. Third, permission to use information is specific to the original purpose. Fourth, the company can keep the data only to satisfy that reason; if the company wants to use the information for another purpose, it needs to initiate a new information collection and use process.

It is often the case, however, that agencies under these regulations do not strictly adhere to the principles relating to the collection and use of personal data. For example, Caudill and Murphy emphasize that adhering to these standards of data collection practices have an antithetical effect on Internet marketing, making it unlikely that online companies, especially those based in the US, will strictly observe them (Caudill and Murphy 2000).

Facebook Data Usage Policy

An SNS usually employs a privacy policy involving how the users' data will be collected by the SNS itself and by third party companies affiliated with the SNS, for the customarily identified aim of "improving user experience". In the past, investigations have found, for example, that Facebook's set-up and privacy policy predisposed it to major flaws that enabled breaches of users' privacy through data mining (Debatin et al. 2009). For instance, Facebook's current data usage policy includes terms and conditions regarding the collection and use of SNS users' data by third party applications. In the past, Facebook allowed third party applications access to almost all user data. Facebook's current data collection policy states that instant personalization applications and other third party websites can be blocked from accessing users' personal information by specifically 'opting out' of allowing the applications access (Facebook 2012). These 'opt out' clauses, however, are often not transparent to the SNS user. It has been reported that before Facebook's set-up was amended in 2007, users' restricted profile information showed up in searches unless the users specifically chose to 'opt out' of their profile from searches (Debatin et al. 2009). Facebook's changing privacy policy reflects the inconsistency between government regulations in regards to the collection and use of personal information for government agencies, and the policies and practices of private online companies.

Implications for Developers

Privacy-Preserving Data Mining

Privacy measures that are inherent in most SNS do not ensure privacy from data mining (Al-Saggaf and Islam 2012). One strategy to prevent the disclosure of personal information from data mining comes from developments within data mining technology itself. Privacy Preserving Data Mining (PPDM) is a development that attempts to prevent a data mining project from misusing personal information (Islam 2008; Brankovic et al. 2007; Clifton et al. 2002). Generally speaking, the aim

of PPDM is to provide the benefits of data mining techniques without allowing a data miner to accurately learn the underlying data being mined, therefore maintaining a degree of privacy in regards to the instances of personal information mined by these techniques (Vaidya and Clifton 2004; Islam and Brankovic 2011). The key to understanding the target of PPDM is to appreciate the difference between so-called "high-level" information (the trends and patterns discovered in a dataset through data mining techniques) and the individual instances of data (the individual pieces of personal information that make up the larger dataset being mined), where the high-level information is what is of value for data miners.

An example of PPDM is noise addition to a dataset in such a way that the data mining patterns of the dataset remain unchanged even after noise addition. However, some (not all) information about some individual is changed. A data miner/intruder does not have any idea which piece of information is changed or which one is not changed, resulting in increased privacy since the data miner/intruder is unsure about the correctness of the observed data of an individual. Note that the technique does not need to add noise to all information, but can still cause great uncertainty regarding the accuracy of observed information. The question that arises is: can developers of SNS take advantage of such a technique by implanting an application that can carefully add noise to some (not all) information of some (not all) users' profiles, if the users choose to opt in for the application? This question is for future research.

One benefit from such an application is that if a data miner sees that a user is interested in both male and female, the data miner will not be sure whether this is the correct information or a noisy one, since there is a chance that some information of some users has changed. To illustrate, using an example from the above findings, while trying to classify a new user as either "lonely" or "connected" if a data miner finds that the new user is interested in both male and female, and does not provide information related to work, then the data miner should normally classify the user as Lonely according to RID 1: Interested In = "Both" and Work = "Absent" → Lonely (33/2). However, due to the noise addition a data miner will have uncertainty about say "Interested In = Both", (as it could originally be "Interested In = Male") and thereby will have less confidence in classification (Islam and Brankovic 2011). Note that, a user needs to opt into use such an application, and the application adds only a little amount of noise to a minimum amount of information for some (not all) users. That way, the application will not change all information about all users, which may make an SNS useless. However, the downside of the technique is that since noise is added only to a low amount of information a data miner can ignore its influence and can guess that the observed information is correct. Thus the data miner can still classify (with some uncertainty) a new user and breach individual privacy.

Implications for SNS Users

While the above two sections discussed implications for governments, business, and SNS developers, users should also do their part to protect their own privacy. Users should mask their information in SNS or add noise to them, which is a commonly known approach among the data mining community (Islam and Brankovic 2011; Brankovic et al. 2007). This approach allows users to carefully disclose some

misleading and false personal information in order to confuse a data miner. If all, or at least most, SNS users take the same approach of placing some false information then a data miner will know that the patterns extracted from the dataset on SNS users are possibly inaccurate. Therefore, it will discourage a data miner from building a dataset on SNS users and extract patterns from such a dataset in order to classify a user as belonging to a group such as "lonely" or "connected". The data miner will know that even if an SNS user, whom the data miner knows, can be classified (through the use of a decision tree or other classifiers) into a group such as "lonely", the classification is probably inaccurate.

Conclusion

This paper explored the potential of data mining as a technique that could be used by malicious data miners to threaten the privacy of SNS users. Using a real dataset, a data mining algorithm was applied to this dataset to demonstrate the ease with which characteristics about the SNS users can be discovered and used in a way that could invade their privacy. It should be noted that there is no difference between anonymised data records and those that contain real persons' names. That is, even if records do not contain real persons' names, data mining can still be used to threaten these persons' privacy.

While SNS users do care about their privacy (Boyd and Ellison 2007; Al-Saggaf 2011), their preference for disclosure (Edwards and Brown 2009) makes them disregard the available privacy measures (Gross and Acquisti 2005; Young and Quan-Hasse 2009). It is essential for SNS users to be careful in maintaining their privacy online. Therefore, it is important to raise their awareness about the possible privacy invasions and their implications. One way to preserve privacy online is by masking the individual's data carefully or hiding the sensitive information such as date of birth, address and other identifying information (Islam and Brankovic 2011; Brankovic et al. 2007). This way, even if an unknown malicious data miner can classify a user, it can be difficult for the miner to locate the user and thus harass him/her. Another way to protect privacy online is by putting restrictions, possibly in the form of laws and regulations, on the use of data mining for individual purposes (Al-Saggaf 2012). However, restricting malicious data mining by introducing laws can be a difficult job. First it will be difficult to detect that someone is mining data maliciously. Second, even if detection is possible, it will be difficult to prove malicious intent. One can always argue that he/she was performing data mining for good intentions such as research and knowledge discovery.

We therefore recommend that the data mining community develop Privacy Preserving Data Mining (PPDM) techniques specifically catered for online environments (Islam 2008, Brankovic et al. 2007; Clifton et al. 2002). Until such techniques are developed, users should protect themselves by using all possible measures, including hiding their identifying information and masking their online activities to protect them from being identified as potential victims (Islam and Brankovic 2011; Brankovic et al. 2007), as well as using SNS privacy settings carefully (Al-Saggaf and Islam 2012). Ensuring the privacy settings on Facebook

are up to date (often they roll back to the default settings) can be another way to maintain privacy in SNS (Al-Saggaf and Islam 2012).

The application of a data mining algorithm on a natural dataset, as demonstrated in this article, has revealed that the threats from data mining on individuals' privacy are serious. In light of these threats, this article has discussed implications for governments, businesses, developers and the SNS users themselves. The major contribution of this article, however, is the use of the decision forest data mining algorithm (SysFor) to the context of SNS to demonstrate, using SysFor itself, how hidden relationships between attributes in a natural dataset can be discovered and used in a way that could invade users' privacy. The decision forest algorithm (SysFor), which was only recently proposed, does not only build a decision tree but rather a forest (i.e. a set of decision trees), allowing the exploration of more logic rules from a dataset compared to the number of logic rules that can be explored by a single tree. In this article the data mining algorithm (SysFor) was applied to the results of a content analysis of 616 Facebook user's profiles which resulted in a natural dataset that contained information on 45 attributes related to those profiles. SysFor built around 800 generated logic rules that were then carefully analysed to identify those strong rules that may be used to compromise the privacy of a Facebook user. It is hoped this innovative approach will stimulate debate among methodologists interested in privacy in SNS.

Appendix

```
Tree 1:
Level (0) Interested-In
  Level (1) Work when b
    Level (2) {c;2,l;31}  when a
    Level (2) Language when p
      Level (3) {c;0,l;2}  when e
      Level (3) {c;0,l;13}  when e+
      Level (3) {c;18,l;10}  when null
      Level (3) {c;0,l;1}  when o
      Level (3) {c;0,l;1}  when o+
  Level (1) {c;62,l;119}  when m
  Level (1) {c;218,l;118}  when null
  Level (1) Religion when w
    Level (2) Family-on-FB when a
      Level (3) {c;0,l;5}  when a
      Level (3) Activities when p
        Level (4) Year-Grad-UC when a
          Level (5) {c;7,l;4}  when a
          Level (5) {c;0,l;1}  when p
        Level (4) {c;0,l;3}  when p
    Level (2) {c;1,l;0}  when p
```

```
Tree 2:
Level (0) Activities
  Level (1) {c;243,1;193}  when a
  Level (1) Timeline when p
    Level (2) {c;64,1;115}  when a
    Level (2) {c;1,1;0}  when p


Tree 3:
Level (0) Language
  Level (1) Timeline when e
    Level (2) {c;1,1;0}  when a
    Level (2) High-School when p
      Level (3) {c;7,1;3}  when a
      Level (3) Year-Grad-UC when p
        Level (4) {c;4,1;17}  when a
        Level (4) {c;7,1;3}  when p
  Level (1) Political-View when e+
    Level (2) Activities when a
      Level (3) Show-Sex when a
        Level (4) {c;3,1;0}  when n
        Level (4) Interested-In when y
          Level (5) {c;0,1;10}  when b
          Level (5) {c;3,1;15}  when m
          Level (5) {c;15,1;7}  when null
          Level (5) University-College when w
            Level (6) {c;4,1;0}  when a
            Level (6) {c;0,1;1}  when p
      Level (3) {c;3,1;41}  when p
    Level (2) {c;1,1;0}  when p
  Level (1) {c;249,1;202}  when null
  Level (1) {c;7,1;4}  when o
  Level (1) Year-Grad-UC when o+
    Level (2) {c;0,1;5}  when a
    Level (2) {c;4,1;0}  when p
```

```
Tree 4:
Level (0) Profile-Image
  Level (1) Timeline when 1
    Level (2) {c;25,1;84}  when a
    Level (2) {c;46,1;72}  when p
  Level (1) {c;27,1;20}  when 10
  Level (1) {c;2,1;0}  when 11
  Level (1) {c;3,1;0}  when 12
  Level (1) {c;0,1;1}  when 15
  Level (1) {c;34,1;45}  when 2
  Level (1) Year-Grad-HS when 3
    Level (2) {c;0,1;1}  when a
    Level (2) {c;2,1;0}  when p
  Level (1) Political-View when 4
    Level (2) Religion when a
      Level (3) Work when a
        Level (4) {c;21,1;0}  when a
        Level (4) Activities when p
          Level (5) {c;10,1;3}  when a
          Level (5) {c;0,1;1}  when p
      Level (3) {c;0,1;1}  when p
    Level (2) {c;1,1;1}  when p
  Level (1) {c;90,1;79}  when 5
  Level (1) {c;1,1;0}  when 6
  Level (1) {c;22,1;0}  when 7
  Level (1) {c;24,1;0}  when 9
Tree 5:
Level (0) Interested-In
  Level (1) DOB when b
    Level (2) {c;0,1;7}  when 1
    Level (2) {c;0,1;1}  when 2
    Level (2) Work when 3
      Level (3) {c;2,1;25}  when a
      Level (3) Language when p
        Level (4) {c;0,1;2}  when e
        Level (4) {c;0,1;13}  when e+
        Level (4) Hometown when null
          Level (5) {c;1,1;3}  when a
          Level (5) High-School when p
            Level (6) {c;0,1;2}  when a
```

```
                Level (6) Year-Grad-UC when p
                  Level (7) {c;12,1;0}  when a
                  Level (7) Family-on-FB when p
                    Level (8) {c;0,1;1}  when a
                    Level (8) {c;5,1;2}  when p
          Level (4) {c;0,1;1}  when o
          Level (4) {c;0,1;1}  when o+
    Level (1) {c;62,1;119}  when m
    Level (1) {c;218,1;118}  when null
    Level (1) Family-on-FB when w
      Level (2) {c;0,1;5}  when a
      Level (2) Activities when p
        Level (3) Year-Grad-UC when a
          Level (4) {c;8,1;4}  when a
          Level (4) {c;0,1;1}  when p
        Level (3) {c;0,1;3}  when p
```

References

Alim, S., Abdulrahman, R., Neagu, D., & Ridley, M. (2011). Online social network profile data extraction for vulnerability analysis. *International Journal of Internet Technology and Secured Transactions, 3*, 194–209.

Al-Saggaf, Y. (2011). Saudi females on Facebook: An ethnographic study. *International Journal of Emerging Technologies and Society, 9*(1), 1–19.

Al-Saggaf, Y. (2012). The mining of data retrieved from the eHealth record system should be governed. *Information Age, 2012*, 46–47.

Al-Saggaf, Y., & Islam, Z. (2012). Privacy in social network sites (SNS): The threats from data mining. *Ethical Space: The International Journal of Communication Ethics, 9*(4), 32–40.

Al-Saggaf, Y., & Nielsen, S. (2014). Self-disclosure on Facebook among female users and its relationship to feelings of loneliness. *Computers in Human Behavior, 36*(2014), 460–468. http://dx.doi.org/10.1016/j.chb.2014.04.014.

BBC News. (2007). Facebook opens profiles to public. Retrieved 12 January, 2012, from http://news.bbc.co.uk/go/pr/fr/-/2/hi/technology/6980454.stm.

Birrer, F. A. J. (2005). Data mining to combat terrorism and the roots of privacy concerns. *Ethics and Information Technology, 7*, 211–220.

Bonneau, J., Anderson, J., & Danezis, G. (2009). Prying data out of a social network. *International Conference on Advances in Social Network Analysis and Mining, 20–22*, 249–254.

Boyd, D. M., & Ellison, N. B. (2007). Social network sites: Definition, history, and scholarship. *Journal of Computer-Mediated Communication, 13*, 210–230.

Brankovic, L., Islam, M. Z., & Giggins, H. (2007). Privacy-preserving data mining. In M. Petkovic, & W. Jonker (Eds.), *Security, privacy and trust in modern data management*. Springer, ISBN: 978-3-540-69860-9, Chapter 11, pp. 151–166.

Catanese, S. A., Meo, D. E., Ferrara, E., Fiumara, G., & Provetti, A. (2011). Crawling Facebook for social network analysis purposes. In *Proceedings of the international conference on web intelligence, mining and semantics*, May 25–27.

Caudill, E. M., & Murphy, P. E. (2000). Consumer online privacy: Legal and ethical issues. *Journal of Public Policy and Marketing, 19*(1), 7–19.

Clifton, C., Kantarcioglu, M., Vaidya, J., & Zhu, M. Y. (2002). Tools for privacy preserving distributed data mining. *ACM SIGKDD Explorations Newsletter, 4*(2), 28–34.

Debatin, B., Lovejoy, J. P., Horn, A., & Hughes, B. N. (2009). Facebook and online privacy: Attitudes, behaviors, and unintended consequences. *Journal of Computer-Mediated Communication, 15*, 83–108.

Edwards, L., & Brown, I. (2009). Data control and social networking: Irreconcilable ideas? In A. Matwyshyn (Ed.), *Harboring data: Information security, law and the corporation*. Stanford: Stanford University Press.

Facebook. (2012). One Billion People on Facebook. http://newsroom.fb.com/News/One-Billion-People-on-Facebook-1c9.aspx. Accessed on October 13, 2012.

Felt, A., & Evans, D. (2008). Privacy protection for social networking platforms. *Workshop on Web 2.0 Security and Privacy*, May 22, pp. 1–8.

Gross, R., & Acquisti, A. (2005). Information revelation and privacy in online social networks. In *Proceedings of the 2005 ACM workshop on privacy in the electronic society*, pp. 71–80.

Harfoush, R. (2011). Has Facebook gone too far? The Mark News Online, November 11. Retrieved January 12, 2012, from http://ca.news.yahoo.com/know-facebook-050204096.html.

Hildebrandt, M. (2009). Who is profiling who? Invisible visibility. In S. Gutwirth, Y. Poullet, P. de Hert, C. de Terwangne, & S. Nouwt (Eds.), *Reinventing data protection?* (pp. 239–252). Berlin: Springer.

Islam, M. Z. (2008). Privacy preservation in data mining through noise addition. PhD thesis in Computer Science, School of Electrical Engineering and Computer Science, The University of Newcastle, Australia.

Islam, M. Z. (2012). EXPLORE: A novel decision tree classification algorithm. In L. M. MacKinnon (Ed.), *Data security and security data*. Berlin/Heidelberg: Springer. LNCS Vol. 6121, ISBN 978-3-642-25703-2, pp. 55–71.

Islam, M. Z., & Brankovic, L. (2011). Privacy preserving data mining: A noise addition framework using a novel clustering technique. *Knowledge-Based Systems, 24*(8), ISBN 0950-7051, (December 2011), 1214–1223.

Islam, M. Z., & Giggins, H. (2011). Knowledge discovery through SysFor: A systematically developed forest of multiple decision trees. In *Proceedings of the ninth australasian data mining conference* (AusDM 11), Ballarat, Australia. December 01–December 02, 2011. CRPIT, 121. P. Vamplew, A. Stranieri, K.-L. Ong, P. Christen, & P. J. Kennedy (Eds.), ACS, pp. 205–210.

Jagatic, T. N., Johnson, N. A., Jakobsson, M., & Menczer, F. (2007). Social phishing. *Communications - ACM, 50*, 94–100.

Johnson, B. (2009). Danah boyd: 'People looked at me like I was an alien'. *guardian.co.uk, at* http://www.guardian.co.uk/technology/2010/jan/11/facebook-privacy. Accessed May 30, 2012.

Johnson, B. (2010). Privacy no longer a social norm, says Facebook founder. *guardian.co.uk, at* http://www.guardian.co.uk/technology/2010/jan/11/facebook-privacy. Accessed May 14, 2012.

Khan, M. A., Islam, M. Z., & Hafeez, M. (2011). Irrigation water demand forecasting—A data preprocessing and data mining approach based on spatiotemporal data. In *Proceedings of the ninth Australasian data mining conference* (AusDM 11), Ballarat, Australia. December 01–December 02, 2011, CRPIT, 121. Vamplew, P., Stranieri, A., Ong, K.-L., Christen, P. and Kennedy, P. J. Eds., ACS, pp. 183–194.

Kirkpatrick, M. (2010). Facebook's Zuckerberg says the age of privacy is over. *ReadWriteWeb*, at http://www.readwriteweb.com/archives/facebooks_zuckerberg_says_the_age_of_privacy_is_ov.php. Accessed May 14, 2012.

Kosala, R., & Blockeel, H. (2000). Web mining research: A survey. *SIGKDD Explorations, 2*, 1–15.

Krill, P. (2011). Big Data mining: Who owns your social network data? InfoWorld.com, March 9. Retrieved 19 December 2011 from http://www.infoworld.com/d/business-intelligence/big-data-mining-who-owns-your-social-network-data-746.

Laurent, W. (2011). The realities of social media data mining. Dashboard Insight, March 14. Retrieved 19 December 2011 from http://www.dashboardinsight.com/articles/new-concepts-in-business-intelligence/the-realities-of-social-media-data-mining.aspx.

Manjoo, F. (2007). Facebook finally lets users turn off privacy-invading ads. Salon.com, December 7. Retrieved 18 December 2011 from http://www.salon.com/2007/12/06/facebook_beacon_2/.

Moor, J. (1990). The ethics of privacy protection. *Library Trends, 39*, 69–82.

Moor, J. (1997). Towards a theory of privacy in the information age. *Computers and Society, 27*, 27–32.

Nakashima, E. (2007). Feeling betrayed, Facebook users force site to honor their privacy. *The Washington Post*, November 30. Retrieved 18 December 2011 from http://www.washingtonpost.com/wp-dyn/content/article/2007/11/29/AR2007112902503.html.

Nissenbaum, H. (1997). Toward an approach to privacy in public: Challenges of information technology. *Ethics and Behavior, 7*, 207–220.

Nissenbaum, H. (1998). Protecting privacy in an information age: The problem of privacy in public. *Law and Philosophy, 17*, 559–596.

Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review, 79*, 119–158.
Nissenbaum, H. (2010). *Privacy in context*. Stanford, CA: Stanford University Press.
Nosko, A., Wood, E., & Molema, S. (2010). All about me: Disclosure in online social networking profiles: The case of Facebook. *Computers in Human Behavior, 26*(2010), 406–418.
Oboler, A., Welsh, K., & Cruz, L. (2012). The danger of big data: Social media as computational social science. *First Monday*, Volume 17, Number 7. Retrieved 24 December 2012 from http://www.firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/3993/3269.
PPIP Act. (1998). NSW privacy and personal information act. http://www.legislation.nsw.gov.au/maintop/view/inforce/act+133+1998+cd+0+N. Accessed 9 June 2014.
Quinlan, J. R. (1993). *C4.5: Programs for machine learning*. San Mateo, CA: Morgan Kaufmann Publishers.
Rachels, J. (1975). Why privacy is important. *Philosophy & Public Affairs, 4*, 323–333.
Rahman, M. A., & Islam, M. Z. (2011). Seed-detective: A novel clustering technique using high quality seed for K-means on categorical and numerical attributes. In *Proceedings of the ninth Australasian data mining conference* (AusDM 11), Ballarat, Australia. December 01–December 02, 2011, CRPIT, 121. Vamplew, P., Stranieri, A., Ong, K.-L., Christen, P. & Kennedy, P. J. Eds., ACS, pp. 211–220.
Rahman, M. G., Islam, M. Z., Bossomaier, T., & Gao, J. (2011). CAIRAD: A novel technique for incorrect records and attribute-values detection. In *Proceedings of IEEE international joint conference on neural networks* (IJCNN 12), Brisbane, Australia. June 10–June 15, 2012, pp. 1–10.
Rubenstein, I. S., Lee, R. D., & Schwartz, P. M. (2008). Data mining and internet profiling: Emerging regulatory and technological approaches. *University Of Chicago Law Review, 75*, 261–286.
Sar, R. K., & Al-Saggaf, Y. (2014). Contextual Integrity's decision heuristic and social network sites tracking. *Ethics and Information Technology, 16*(1), 15–26.
Sar, R. K., Al-Saggaf, Y., & Zia, T. (2012). You are what you type: Privacy in online social networks. In S. Leitch & M. Warren (Eds.) *Proceedings of the Sixth AICE conference*, Melbourne, Australia, 13 February 2012 (pp. 13–18). Deakin: School of Information Systems, Deakin University.
Tavani, H. T. (1999). Informational privacy, data mining, and the Internet. *Ethics and Information Technology, 1*, 137–145.
Tavani, H. T. (2011). *Ethics and technology: Controversies, questions, and strategies for ethical computing* (3rd ed.). Hoboken, NJ: John Wiley.
Thelwall, M., Wilkinson, D., & Uppal, S. (2010). Data mining emotion in social network communication: Gender differences in MySpace. *Journal of the American Society for Information Science and Technology, 61*, 190–199.
Ting, I. (2008). Web mining techniques for on-line social networks analysis. In *International conference on service systems and service Management*, June 30 2008–July 2, pp. 1–5.
Vaidya, J., & Clifton, C. (2004). Privacy-preserving outlier detection. In *Proceedings of the 4th IEEE international conference on data mining* (ICDM 2004), pp. 233–240.
Van den Hoven, J. (2008). Information technology, privacy, and the protection of personal data. In J. van den Hoven & J. Weckert (Eds.), *Information technology and moral philosophy* (pp. 301–321). Cambridge: Cambridge University Press.
Van Wel, L., & Royakkers, L. (2004). Ethical issues in web data mining. *Ethics and Information Technology, 6*, 129–140.
Young, A. L., & Quan-Hasse, A. (2009). Information revelation and internet privacy concerns on social network sites: A case study of Facebook. *Proceedings of the fourth international conference on Communities and technologies, 2009*, 265–274.