

Assisted Social Identity Management

Enhancing Privacy in the Social Web

Michael Netter, Moritz Riesner, Günther Pernul
Department of Information Systems
University of Regensburg
Regensburg, Germany
{firstname.lastname}@wiwi.uni-regensburg.de

ABSTRACT

The rise of the social web shifts personal identity management to the online world. As a result, personal information is persistently available to all of a user's contacts without distinguishing between different contexts such as *Work* and *Leisure*. Personal information being available to audiences outside the intended context violates contextual integrity and poses a threat to users' privacy. We argue that a formal description and a conceptualization of the problem scope is required to systematically address current challenges of personal identity management in online social settings. Based thereupon, we propose assisted social identity management to support the user in finding segregated audiences among his contacts as a first step to advance from the current situation and sketch further improvements. We evaluate our approach using real-world data, demonstrating the feasibility of our proposal.

Keywords

Assisted Social Identity Management, Audience Segregation, Contextual Integrity, Social Web, Privacy

1. INTRODUCTION

Besides the established services of the WWW – distribution of information and commercial use – the social web, consisting of communication and interaction between users enabled through easily usable applications, has gained importance rapidly [19]. Today it offers an infrastructure for communication, information and self-expression. Even established services on the Internet, such as news portals incorporate related features, thus adding to the pervasiveness of the social web. A current example is the “Like-Button” introduced by the 500-million user social networking platform Facebook that can be embedded into regular websites to connect their content with user-generated comments and recommendations. In enterprise settings, social applications generate new opportunities such as tapping into new knowledge pools originating from employees connected by social

networks and assessing potential for innovation through better customer integration into the value chain [2].

This rise of relevance and quantity of social web services has also been accompanied by concerns regarding privacy and thus raises organizational and legislative implications. Politicians criticize possible data protection issues, for example with services offered by Google and Facebook, and demand stricter regulation. Similarly, users of the social web voice demand for more protection of privacy as well.

While this common criticism usually targets enterprises and their handling of customer data, there is also a need for research regarding the sharing of personal data between users of the social web. In fact, as the social web is used to shift personal identity management (IdM) to the online world, the differences to the offline world pose new challenges. In current implementations of the social web, personal information is available persistently, digitally and thus unchanged. Contrary, in the real world, personal information is transient and availability is limited to certain settings. Existing notions of privacy, usually involving confidentiality [18] and access control [20][5], cannot be applied without modification, as users of the social web have the desire to selectively share personal information depending on situation and recipient. However this demand for personal information on the one hand being publicly available, but on the other hand not for the public [16] poses a challenge to users of the social web. This trade-off between privacy and comfortable disclosure of personal information is referred to as “privacy paradox” [13], calling for contextually aware information sharing. Unlike the real-world, where different contexts such as coworkers and a group of friends are automatically separated by mere physical and spatial separation, providers of services for the social web do not distinguish between different contexts, thus often exposing the complete digital identity of a user to her contacts.

Identity management is commonly used in organizations and enterprises to administrate individuals and control their access to resources. On the contrary, social IdM is a user-centric concept to cope with the challenges of presenting different facets of the self to different audiences and to keep those views consistent. In this work, we introduce ways to improve social IdM by distinguishing between different contexts. This enables context-aware segregation of audiences of the online identity which in turn can be realized by the possibility to present different values for the same attribute

to different audiences. The proposal is accompanied by a formal presentation of the problem scope and by a prototype application aiding the user in managing her online contexts.

The remainder of the paper is structured as follows: In the next Section, we give an overview of related work in the field of social IdM and audience segregation. In Section 3, we show that contextually segregated audiences can be realized through directed attribute presentation and introduce the problem scope in a formal fashion. After conceptualizing the problem scope we propose several approaches to implement assisted social IdM that support the user in recognizing and organizing online contexts in Section 4. The prototypical implementation of two of these approaches and a subsequent evaluation are presented in the following Section 5. Finally, in Section 6 we summarize our findings.

2. RELATED WORK

The increasing relevance of the online digital identity has been recognized by a plethora of research. Nowadays, personal data resides in a growing number of databases, while information on the Internet becomes easier to find and harder to hide [11]. The users' digital identity is defined as their attributes and their values [4], while a subset of them constitutes a partial identity as described in [14] and [8]. Cameron sees the increasing number of personal digital identities as a problem, as users lose track of their disclosed data [3]. While presenting different attribute values to different audiences is related to the concept of partial identities and these notions are related to our work, these works are mostly concerned with interaction with service providers, while we focus on selected disclosure of personal data to contacts in the social web.

This difference can also be seen when looking at the state of end-user IdM systems (IMS), which focus mostly on account management issues such as Single-Sign-On. The need for IMS offering "context-dependent role and pseudonym management" has been postulated [1], however, since the identity of a contact in the social web is usually disclosed, we see a shift from pseudonym management to context-based attribute management.

The notion of privacy as contextual integrity was first introduced by Nissenbaum [12] to cope with the challenges of information technologies processing personal data. She argues that contextual integrity is a measure of how closely gathering and dissemination of personal information conforms with the intended context for this information. Privacy is breached if that personal information is available outside this context. This requires the user to be free from constraints on the construction of his identity as stated in the concept of privacy as practice that is introduced in [7].

To facilitate contextual integrity, audience segregation is a valuable concept. Originally developed by Goffman [6], it states that each individual performs multiple and possibly conflicting roles in everyday life, and it needs to segregate the audiences for each role, in a way that people from one audience cannot witness a role performance, that is intended for another audience and thereby keep a consistent self-image. In [17] this concept is applied to social networks, emphasizing the increased complexity of audience segrega-

tion compared to the everyday life.

A few Facebook applications share similarities with the prototypical implementation presented in Section 5.1, which creates segregated audiences by grouping the user's contacts based on their mutual friends. These applications differ from our work by aiming at a visually appealing presentation of a user's network of contacts while this work focuses on finding disjoint segregated audiences in the context of assisted social identity management. Friendwheel¹ allocates all contacts on a circle based on their proximity to each other, which is derived from mutual friend connections. Unlike our approach however, it doesn't create disjunct groups of contacts. TouchGraph² and Social Graph³ do create disjunct contact groups. Yet, for all of these implementations, neither the algorithm, nor the exact input parameters to compute the friend groups are available. Also, these implementations are specifically built for the Facebook platform while we aim for an approach that is agnostic of a particular social network by solely employing contact relations for finding segregated audiences.

The social web is increasingly being recognized by EU research projects. Padgets⁴ acknowledges the relevance of the social web and harnesses its community knowledge to get input for policy making while being committed to preserve participants' privacy. PrimeLife⁵ [15][10], a European research project, employs the concept of audience segregation to develop an advanced social network which allows to define different social contexts and assign different audiences. This work differs from ours as we propose means to support the user in finding segregated audiences within her contacts while PrimeLife focuses on how to assign contacts to different social contexts and thus here audiences must be defined manually.

3. SEGREGATED AUDIENCES IN SOCIAL IDENTITY MANAGEMENT

In this section, we explain the need for audience segregation and directed attribute presentation as a means to implement it in the social web, followed by a formal definition of the problem scope.

This work is agnostic to how the online identity is hosted and how the communication between contacts is organized and implemented. While possibly enabled through a single social networking website, other ways to implement social IdM are also possible, such as distributed solutions. Also, we only consider other members as a possible audience while privacy issues regarding entities such as the operator of a social networking site or other service providers are out of scope of this work. Furthermore, we do not consider possible inferences through publicly available information in connection with shared personal data through linkability.

¹<http://thomas-fletcher.com/friendwheel/>

²<http://www.touchgraph.com/TGFacebookBrowser.html>

³<http://www.mihswat.com/labs/app/facebook-social-graph/>

⁴Policy Gadgets Mashing Underlying Group Knowledge in Web 2.0 Media - <http://www.padgets.eu>

⁵Privacy and Identity Management in Europe for Life - <http://www.primelife.eu>

The research is carried out from the point of view of a person (“user”) interacting with other individuals through an online social setting. Those other individuals are commonly viewed as “contacts” or even called “friends”. To construct an attacker model, one has to see them as privacy-attacking adversaries that may be able to discover information that was meant to be kept undisclosed to them. We define the attacker as “honest but curious”, acting within the rules set by the services providing the social web, thus only accessing legitimately available information and not exploiting security weaknesses of the site or, for example, weak passwords. Therefore, the technical protection of the confidentiality of certain attribute values is out of scope of this work, while we focus on the assignment of attribute value visibility to contacts depending on contexts.

3.1 The Need for Segregated Audiences

Successful personal IdM in the social web requires means to organize one’s identity depending on social contexts and to act accordingly. The user needs to be supported in recognizing and distinguishing between different social contexts online similarly to offline contexts such as work, school, family and friends. Like with offline contexts, which each have a different appropriate and accepted behavior, one must be able to choose how to present one’s identity depending on online contexts as well. As one performs multiple and possibly conflicting roles in various contexts, the need for audience segregation occurs, meaning that different audiences are kept from witnessing role performances that were meant for other audiences [6].

An identity in the social web consists of several attributes and their associated values, while individuals being able to view certain attribute values are referred to as an audience. In the online social web, *directed attribute presentation* is needed to realize audience segregation. In other words, one must be able to present different values for the same attribute to different audiences and to hide certain attribute values from other audiences, thus keeping presented partial identities consistent. One example that has been brought up is a teacher that feels the need to hide certain spare time activities and friends from her student contacts on a social networking site after irritations about her hobby occurred [15].

As in real life, audiences need to be disjoint with each contact only seeing the predetermined attribute values and thereby preserving the integrity of the user’s partial identity. In case of overlapping audiences, i.e. contacts spanning several segregated audiences, these contacts are assigned to a newly created audience. Enforcing this policy raises the problem of two of the users’ contacts from different segregated audiences possibly exchanging information on the user without his knowledge and thereby violating the contextual integrity, which is unsolved even in the real world.

A context, describing each instance in which certain attribute values and a particular audience come together [15], can be seen as a means to connect sets of attribute values and their corresponding audiences. Thus, users must be supported – assisted or automatically – in allocating both online contacts and appropriate attribute-value pairs, thereby forming a specific context. Figure 1 shows two sample con-




Figure 1: Two Sample Contexts of User Bob

texts, one for work and one for spare time activities. Some attributes, such as the salary, are only shown in one context, some, such as the first name have different values in each context while others are shown in both contexts.

Presently however, directed attribute presentation is not or only in a limited way implemented in the field of the social web. We argue that such controls are essential to empower users for context-based identity presentation. Currently, users wishing to customize their digital appearance to different audiences need to resort to workarounds such as creating multiple accounts at the same or distinct social platforms. By doing so, they are showing different sides of their identity to different audiences and thereby create distinct contexts. The option to hide certain attributes from a subset of one’s contacts is already available on some sites, but does not reach far enough.

There are several reasons for the current lack of more fine-grained customization possibilities, for example providers of social platforms may not have seen a need to implement such features due to a lack of customer interest and business value. Also, the complexity of such controls poses a challenge both for developers and users. We improve the situation by showing further approaches to segregate audiences and to enable directed attribute presentation in Section 4.

3.2 A Formal Problem Scope Description

Followingly, we express the notions related to contexts, audiences and directed attribute presentation formally and thus more precisely to ensure a clear presentation of the problem scope and to support future work. Firstly, for the user, there is a set of attributes

$$A = \{a_1, a_2, \dots, a_n\}.$$

For each of these attributes, the user may enter a set of values

$$AV_m = \{av_{m1}, av_{m2}, \dots, av_{mk}, \emptyset\} \forall m = 1 \dots n.$$

Always included is the empty value representing that the attribute is not shown. Note that in current services the user is restricted to only one attribute value and the empty value. Further, there is a set of contacts


$$C = \{c_1, c_2, \dots, c_p\}$$

denoting user’s contacts in the social web.

For the selected, directed and context-based release of attribute values, the following question must be answered: Which value av_x of the attribute a_x is presented to contact c at the time of access? To answer this question, one


3

Automated Social IdM



2

Assisted Social IdM



1

Manual Social IdM




Figure 2: Conceptualization of Current and Proposed State of Social Identity Management

also needs to consider on which information, denoted as i , the decision should be based, which can be summarized by a function

$$f(a, c, i) \Rightarrow av$$

f can be seen as a *view* on an attribute and may be invoked each time a contact attempts to access an attribute of the user. Similarly, we define a view V on all of the user's attributes as the set of attribute values

$$F(c, i) \Rightarrow V \subseteq AV_q \forall q = 1..n.$$

It returns the set of attribute values a particular contact may see, thus representing a partial identity of the user.

Before elaborating further on the possible ways to implement f , we need to consider the previously mentioned audience segregation. We consider a segregated Audience (SeA) as a subset of the user's contacts C , therefore

$$SeA_i = \{c_1, \dots, c_r\} \subseteq C.$$

As the purpose of SeAs is to prevent the presentation of inconsistent partial identities or single attribute values, all contacts that form a SeA have to be presented with the same attribute values leading to

$$f(a_x, c_y, i_y) = f(a_x, c_z, i_z) \forall x \subseteq 1..n; c_y, c_z \subseteq SeA.$$

In other words, f has to return the same attribute values for all contacts in the same SeA. In order to prevent contacts from seeing different and therefore conflicting attribute values, one contact can only be part of one SeA, thus all SeAs are disjoint subsets of C . To resolve the case of overlapping SeAs, they may be split up so that the intersection forms a new SeA.

We define a context as a set of contacts forming a SeA and a corresponding consistent set of attribute values that

may be exposed to these contacts, as illustrated in Figure 1. Thus, F , returning an attribute value set for a given contact can be used to establish the connection between contacts and attributes that is necessary to define a context. To reduce complexity, the problem of defining F can be split into two tasks represented by Functions G and H , namely assigning contacts to contexts and the corresponding SeA ($G(c, j) \Rightarrow SeA$) on the one hand and assigning attribute values to audiences ($H(SeA, l) \Rightarrow V$) on the other hand with j and l being additional information on which the decision is based on.

One straightforward way to implement G and H are manual assignments conducted by the user, thus resulting in j and l being static lists. However, with the complexity arising from the ever-increasing number of contacts, the abundance of attributes and the added possibility of using multiple values for each attribute, this task is prone to become tedious at best. Further, the dynamic nature of the social web would not be considered. Contacts previously not classified, changes of contact's attributes and the nature of the relationship with them are partly out of control of the user and would require reconsidering the attribute values released to them. Also, notions such as SeAs, directed attribute release, contexts and views may not be intuitive to the average user. These challenges require an advancement from the current state of the art both conceptually and as seen in current social web implementations and are addressed in the following sections.

4. APPROACHES TO DETERMINE SEGREGATED AUDIENCES

In this section, we propose a conceptualization of social IdM, showing that currently users have to manually deal with the complexity imposed by the social web. We advance the situation by proposing assisted social IdM and illustrate several

Table 1: Proposed Audience Segregation Approaches

No.	Type	Description	Required input data j	Categorize new user at time of access	Combinable with
1	Manual	Manual group creation and user assignment	Manual SeAs definition	no	2, 3, 4, 5
2	Attribute Rules	Dynamic assignment based on predefined attribute rules	Set of Rules	yes	1
3	Contact Interaction	Analyze past user behavior and interaction	Quantity and quality of past contact interaction	yes	1, 5
4	Attribute similarity	Assign new users according to similarity to predefined and populated classes	Manual SeAs definition, contact’s attributes	yes	1, 5
5	Contact Relationships	Determine distinct contact groups	All contacts’ contact lists	yes	1, 3, 4

approaches to support the user.

4.1 Current and Proposed State of Social IdM

Today’s social network landscape consists of many service providers offering platforms for different contexts such as leisure (e.g. Facebook) and work (e.g. LinkedIn). However, the possibilities to cope with multiple contexts within the same platform are limited, as available possibilities to enforce audience segregation and ways to exert directed attribute presentation leave much to be desired. This is shown in Figure 2, in which the current state is illustrated at the bottom level as *Manual Social Identity Management*. While it is already possible for the user to assign contacts to groups and set attribute visibility based on group membership, all of these steps have to be performed manually by the user before the attribute in question is accessed by a contact. It is currently not possible to show different attribute values of the same attribute to different contacts. Additionally, current provider implementations lack default groups with settings that provide only limited access to personal information. A further shortcoming is the lack of transparency for the user about who has access to which disclosed information stemming from large contact group sizes and inappropriate tools to visualize audiences. Thus it is difficult to act appropriately for a given context and thereby maintain contextual integrity. As mentioned in Section 3.1, the common approach to cope with this limited status quo is to employ workarounds.

We propose improving the current state by managing segregated audiences through directed attribute release, e.g. displaying different attribute values to contacts depending on the current context. As the transition to the use of different attribute values increases the complexity of social IdM even more, the need for *Assisted Social Identity Management* arises, e.g. the user needs to be supported in creating and managing groups of contacts that correspond to SeAs and their respective contexts. As illustrated in the middle level of Figure 2, this assistance encompasses automatically generating suggestions for SeAs, namely allocations of the user’s contacts to groups. Thus, the initial burden of creating groups from possibly hundreds of contacts is lifted. The user may then refine these suggestions manually, approve them and assign attribute values to form a context. In Section 4.2, we present and discuss four new approaches to generate SeA suggestions.

While assisted social IdM eases the initial allocation of existing contacts into groups, there is still room for improvement: Firstly, any suggested group allocation needs to be verified manually and refined if necessary, as errors would possibly lead to the unintended release of personal information. Also, after user approval, no further information will be considered for the access definition. Yet, dynamic incorporation of new information would be desirable, for example to prevent a coworker that has left the company from viewing work-related attributes. Thus, we propose *Automated Social Identity Management* as a next step to improve social IdM, as shown on the top level in Figure 2. It advances from assisted social IdM through the following properties:

- 1. User parametrization of contact allocation:**
Rather than having to decide or approve each contact’s group individually, the user sets parameters for automatic contact allocation.
- 2. Contact allocation after user interaction:**
Unlike as in assisted social IdM, user interaction occurs before contacts are assigned to groups, thus allowing for dynamic contact-to-audience allocation and incorporation of the most recent available information. However, this poses high requirements on the algorithm’s reliability.

Besides their adequacy to produce suggestions for SeA, in Section 4.2, the aforementioned approaches are examined regarding their use in a possible automated social IdM scenario.

4.2 Approaches to Segregate Audiences

In the following we develop and discuss more advanced approaches to support the user in creating and managing segregated audiences in order to maintain contextual integrity. The goal of all approaches is to create disjoint groups of contacts, or, applying the formal notation introduced in Section 3.2, to define parameter j of function $G(c, j) \Rightarrow SeA$. Table 1 provides an overview of our proposed approaches.

Manual group assignment. This static approach requires the user to manually cluster her contacts in a reasonable manner and thereby create segregated audiences. Using

manual assignment, j simply represents the user’s knowledge on each contact, i.e. she has contextual information about the contact and is therefore able to assign it to an appropriate segregated audience.

Theoretically, this approach allows for a fine-grained assignment of all contacts. However, with the increasing number of contacts, the user has to deal with the complexity of overlapping audiences which need to be split up, in order to ensure consistency of the presented attribute values. Furthermore, due to the static nature of this approach the user has to reconsider the classification each time a new contact is added. In summary, manual assignment is a powerful concept to create and manage segregated audiences, however due to its high complexity it needs to be combined with other approaches that aim at supporting the user. This approach is implemented in many of today’s social networks and is listed for reference purposes.

Attribute rules. Employing attribute rules makes use of a contact’s attributes to assign her to a segregated audience. In more detail, the user predefines a set of rules that constitute the input j of function G . A simple rule might be $R : \{av_{lastname} = Doe\} \Rightarrow SeA_{Family}$ to assign each contact having *Doe* as the attribute value for the attribute *lastname* to the segregated audience *Family*. More sophisticated rules are conceivable to get finer-grained segregated audiences. A set of predefined rules for reoccurring problems could be provided to the user. As it is possible to create contact groups using this approach can be employed in assisted social IdM to make SeA suggestions based on the user’s preferences.

Once a rule set is defined thoroughly, this approach can operate without further user interaction, hiding all the complexity from the user and allowing for automated social IdM. New contacts can be classified automatically and segregated audiences can be created dynamically at access time. However, the assignment precision of a contact and thereby the effectiveness highly depends on the quality of the rule set. As the assignment decision is based on the contact’s attributes, without further assurance about those attributes (e.g. trust authority), this approach is vulnerable to a malicious contact that alters his attributes in order to be assigned to a different segregated audience and thereby seeing different attributes and attribute values.

Contact Interaction. This approach relies on the user’s preexisting interaction with her contacts and the fact that both quality and quantity of interaction depends on the contact’s context. In other words, we propose to analyze both the frequency and the content of all previous communication that is available (e.g. private messages, wall postings, etc.). As an example, a user might frequently communicate with her best friend using everyday language while she might have written only few formal messages to her colleague.

In a more formal way, qualitative and quantitative communication patterns constitute the input j of function $G(c, j) \Rightarrow SeA$. As a result, function G delivers a set of segregated audiences that can be further refined by the user. New contacts without any prior interaction must be assigned to a




Figure 3: Relationships Between Contacts

default group requiring the user to finally assign them.

Besides the use of this approach for SeA suggestions in assisted social IdM, applying it to automated social IdM is also feasible. For instance, contacts without interactions in a certain time period could be automatically put into a group with less access privileges.

Attribute similarity. This approach is based on the assumption that all people in a given context are similar to some extent or share a common attribute (e.g. the same affiliation). We propose to make use of this similarity between a user’s contacts to create segregated audiences.

Two variations of employing attribute similarity can be conceived: First, in order to find meaningful segregated audiences among all contacts, unsupervised learning techniques, such as clustering algorithms can be applied. Without any user interaction, all contacts are assigned to subsets in a way that all contacts in one subset share similarities. The approach is useful to suggest a set of segregated audiences which the user can adapt and refine and is thus use in assisted social IdM. For example this approach can be used to discover the segregated audience $SeA_{FellowStudents}$ since all associated contacts have the same affiliation and are of similar age.

Second, supervised learning techniques can be used to classify new contacts based on preexisting groups. In more detail, after the user has defined segregated audiences (e.g. using one of the other approaches depicted in Table 1), a classifier is determined and used to suggest the correct segregated audience for any new contact.

For the first variant of the approach the input j of function G is a vector of attributes that are suitable for similarity measurement. The second variant additionally requires the predefined segregated audiences as input value. Machine learning techniques are a powerful concept to cope with the high complexity of finding segregated audiences among a user’s contacts [21]. However as the algorithms lack contextual knowledge the assignment precision is low and therefore the user has to finally decide upon the classification. Similar to the rule-based approach, the decision is made upon a contact’s attribute and is therefore vulnerable to a malicious contact that changes her attributes in order to be classified with another context. Thus, the use in automated social IdM is questionable.

Table 2: Contact Groups from the Test Scenario

No.	Name	Description	Group Size
1	Hometown	Contacts from the student’s hometown	3
2	High school year	Contacts from a student exchange year spent in the United States	8
3	Study abroad year	Mostly international contacts from a study abroad year during college	25
4	USA 2	Contacts in the US unrelated to group “High School Year”	8
5	Summer school	Contacts from an international summer school attended during PhD studies	18
6	University	Contacts known through the University	77
7	Other	Other contacts not related to groups above	6
Sum			145

Contact Relationships. The rationale of this approach is that information about a contact’s context is embedded in the relations this contact has with other contacts. We propose to employ this inherent property of common social networks to arrive at segregated audiences. It can be seen as a special case of the attribute similarity approach with the contacts of a contact as the only attributes being considered.

To be more specific, both the user and a specific contact have relations to other contacts as depicted in Figure 3. It can be seen that the user’s contacts c_1 , c_2 , and c_3 have a relation to each other and thereby form $SeA_1 = \{c_1, c_2, c_3\}$. Analogously the contacts c_4 , c_5 constitute $SeA_2 = \{c_4, c_5\}$.

We assume that contacts that have relationships among each other also communicate with one another and possibly about the user and her attribute values. This approach has the benefit that it makes communication between members of different SeA and therefore the discovery of different attribute values by the same contact less likely. However, this cannot be prevented completely for two reasons: First, not all relationships in real life are also represented online in the social web, therefore there may always be links between contacts that cannot be discovered. Secondly, only allowing SeA allocations in which there are no links between members of different SeAs would be the strictest form of this approach. It would possibly lead to only a few or one big SeA, as there are usually relations between contacts from different contexts such as a user’s friend knowing a user’s colleague from work (note that the two SeAs in Figure 3 satisfy the requirement). In order to find reasonable segregated audiences it is necessary to allow linkage between the contacts of two SeAs.

The approach will gain importance even more once the current service provider landscape with its distinctive social networks for different contexts (e.g. Facebook for *Leisure* and LinkedIn for *Professional*-activities) converges to one large or several interconnected networks for all contexts.

To arrive at segregated audiences, this approach uses information about relations between users and contacts as input parameter j . The approach can be employed to discover segregated audiences within the user’s contacts and enable assisted social IdM, which is demonstrated in Section 5. It is also conceivable to employ such an approach of automated

social IdM by classifying new contacts based on their relationships to contacts in existing groups.

5. IMPLEMENTING ASSISTED GROUP ASSIGNMENT

In this section we evaluate the effectiveness of *Contact Relationships* as one of our proposed approaches to discover segregated audiences among a user’s contacts. We first present our algorithm that implements the function G as introduced in Section 3.2. Subsequently, we describe the test scenario based on a test person’s real-world data and employ this dataset to evaluate our results using a prototype.

5.1 Algorithmic Foundations and Prototypical Implementation

The test person’s friend list and their connections to each other were gathered using an application connecting to the social networking site Facebook. With *FQL (Facebook Query Language)*⁶, the site allows applications to access some of the information visible to their users through a number of available tables. We queried the table *connection* to get a list of the user’s contacts and the table *friend* to obtain the links between those contacts.

We further implemented a prototype to derive segregated audiences within the social networking site dataset and evaluated the effectiveness of our approach. The prototype is available for download and testing on our website⁷. For evaluation purposes it allows for manual creation of segregated audiences (cf. Manual approach in Table 1) that serve as a test set. The tool contains an algorithm implementing the *Contact Relationship* approach to find segregated audiences among a user’s contacts and a matching algorithm to evaluate the result against the predefined test set (see Figure 4). Furthermore we included a simulator to rapidly determine the optimal initialization parameters for the *Contact Relationship* algorithm.

The prototype groups the test person’s contacts into classes, using the relationships among them as a criterion, based on the assumption that contact groups with a high number of relationships among them also belong together semantically

⁶<http://developers.facebook.com/docs/reference/fql/>

⁷<http://www-ifsresearch.wiwi.uni-regensburg.de/paper/wi/>

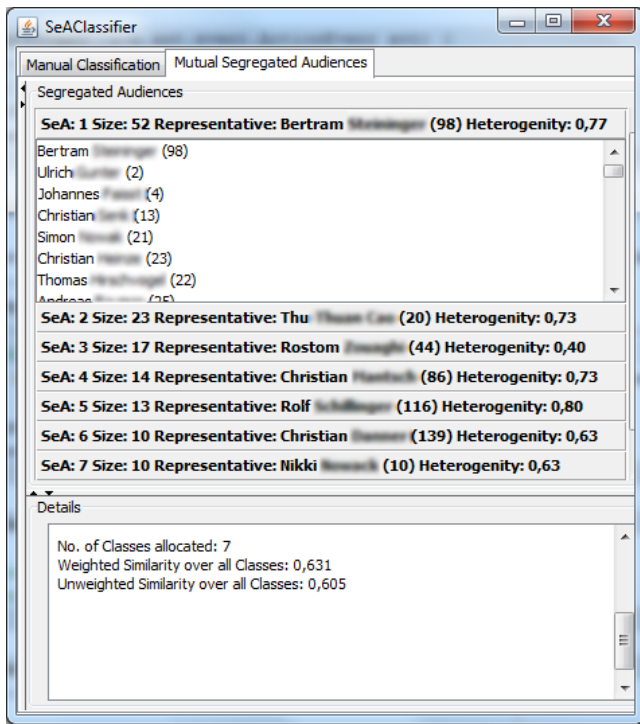


Figure 4: Our Prototype (Screenshot)

and thus form a possible context. A “relationship” is a binary property derived from the “is friend” property that popular social networking sites employ.

For clustering, we employed a heuristic that uses the Jaccard index [9] – a frequently used distance metric – to calculate the similarity between two contacts. First, it calculates a matrix containing similarities between any combination of two contacts based on the overlap of their contact sets employing the Jaccard similarity coefficient. Thus, two contacts with a high number of mutual contacts receive a high similarity value. Next, representatives for the first two classes are selected, with the first one being the contact with the highest number of similarity values that exceed a predefined threshold and the second one being the contact that has the lowest similarity to the first representative. Further class representatives are chosen based on the maximal distance to all existing representatives as a criterion, until the distance of the next possible representative falls below the previously chosen threshold. Next, all remaining contacts are assigned to classes based on their similarity to the class representatives. The threshold greatly influences the outcome of the algorithm, as the number of created classes is directly related to it. In Section 5.3, a detailed discussion on the optimal threshold determination is provided.

In terms of machine learning, this approach is a type of unsupervised classification, as it requires no prior training of the algorithm by the user. Note that the existence of a connection between two contacts does not weigh more towards their similarity value than any other mutual contact they have. Also, only the existence of connections and thus mutual friends influence the outcome, not their absence. This

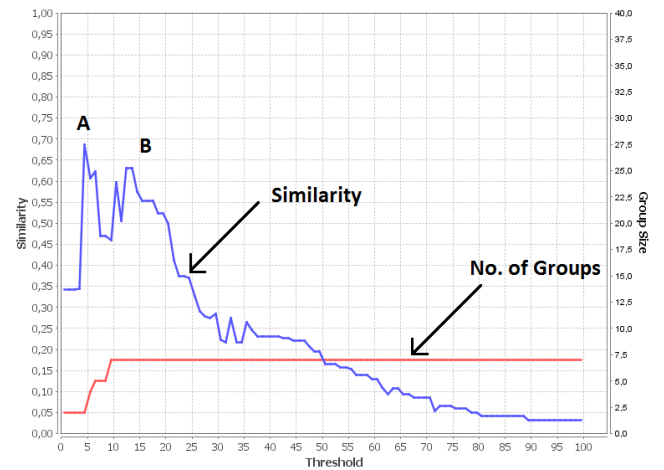


Figure 5: Threshold Determination

matches the nature of the data set, in which instances where two contacts have no connection far outweigh the number of connections.

5.2 Test Scenario

For evaluation, we obtained actual user data containing 145 contacts and their relationships from a PhD-student’s social network account and created SeAs automatically using the algorithm. Then, we compared this output to a manual assignment that the test person was asked to create.

The test person was asked to classify his contacts into semantically meaningful groups similar to real-life contexts to the best of his knowledge, leading to the allocations presented in Table 2. Groups such as the international *summer school* and *study abroad year* containing contacts from various countries and with differing affiliations show that a classification algorithm cannot be realized trivially by only using these attributes for classification.

5.3 Evaluation

We employed a cluster to classes evaluation methodology [21], i.e. the test person was asked to create a set of optimal segregated audiences, as shown in Table 2. The manual classification was then compared to the classes our algorithm created. In more detail, starting with the largest manually created group the most similar automatically created group was chosen for evaluation and a similarity value was determined based on member overlap. If the algorithm created more groups than defined in the manual classification, the smallest manually created group was compared to a backup-group, that was created of all members of the residual automated groups.

As mentioned before, the quality of the results highly depends on the similarity threshold the algorithm uses to create classes of segregated audiences. To determine the optimal threshold we employed a simulator implemented in our prototype, calculating the average cluster to classes similarity for each threshold between 0 and 100 percent. The results for are depicted in Figure 5. As explained in the previous paragraph, if the algorithm finds more classes than

Table 3: Classification Results

No.	SeA_{man}	SeA_{man} Members	SeA_{aut} Members	Sim.
1	University	77	52	0.633
2	Study abroad year	25	23	0.846
3	Summer school	18	17	0.842
4	High school year	8	8	1.000
5	USA 2	8	10	0.800
6	Other	6	14	0.111
7	Hometown	3	23	0.000
Weighted similarity				0.631

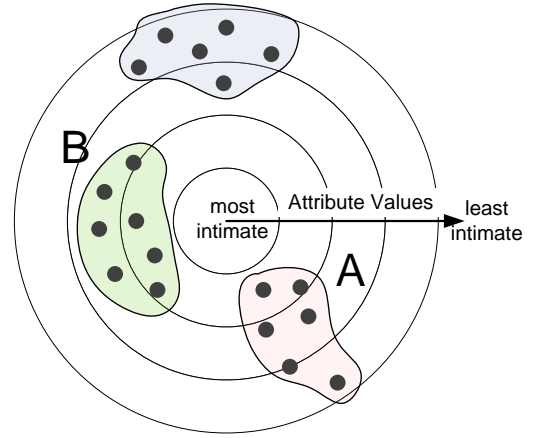
in the test set, a backup class is created. This is shown by the red line, which is capped at the number of user-defined classes. The optimal threshold is not at peak A, as the algorithm only creates two classes with that threshold, whereas at peak B the number of found classes equals the number of classes in the test set and thus peak B denotes the optimal threshold.

Using a threshold of 13.5 percent similarity, Table 3 depicts the classification results for our algorithm and the cluster to classes evaluation against the predefined set (cf. Table 2). The similarity of each pair of classes (manual and automated) is weighted with its average number of members resulting in an overall class similarity of 63.1 percent. This means the automatically created classes overlapped the manually created classes by 63.1 percent in average with the average being weighted using the class sizes.

As Table 3 shows, the result quality highly depends on the contact group. For example, for the classes *Study abroad year* and *Summer school* there is a high similarity between manually and automatically created classes. This can be explained by the high interconnection between the group members and their distance to other groups. As the algorithm highly depends on interconnections between group members, the results are less optimal for other contact groups like *Hometown*. As the members in this group do not know each other, this explains the algorithm’s inability to find an appropriate segregated audience for that group.

The results clearly demonstrate the feasibility of automated finding of segregated audiences within a set of contacts which can be further refined by the user and thereby implement assisted social IdM.

A further evaluation comparing our approach with the related Facebook applications mentioned in Section 2 would be desirable. Yet, due to the different nature of the approaches, one could not rule out misleading findings. Friendwheel does not produce any contact groups, thus providing no results that could be compared. For the other two approaches, due to the limited available documentation, one cannot assess if they consider any other input parameters besides the contact’s relationships among each other. Also, we consider a clear understanding of the employed algorithms as a precondition before comparing outputs in a meaningful way.

**Figure 6: Audience Distribution**

If such a comparative evaluation were carried out, it would have to be performed on exactly the same input data, ideally by adding the other application to the test user’s Facebook-profile. The cluster to classes evaluation could then be applied to compare the resulting classes to the manually created user groups, yielding a similarity value that would be easy to compare.

5.4 Ideas for improvement

While our implementation has shown the feasibility of assisted social IdM, there are various angles on which both the algorithm and the general idea of assisted IdM can be improved.

The presented algorithm produces promising and usable results, yet there are various other clustering approaches [21] with a number of possible settings and parameters which allow for further research on their suitability in this problem setting. Furthermore, using other input attributes in addition to the contact’s mutual friends opens up future work opportunities to further optimize our approach. Still, groups such as *Summer School* (cf. Table 2) show that members of valid segregated audiences may have little or nothing in common besides mutual friends.

The focus of the demonstrator lays on the assignment of contacts to groups, corresponding to function G described in Section 3.2. For a full specification of contexts in social IdM however, an assignment of attributes to contexts is also necessary. The attribute-to-context assignment could happen after the assignment of users to contexts, however there could be interdependencies. For example, adding a certain attribute visibility to a context may require reconsidering the assigned contacts and splitting the group, as only some of the corresponding contacts were meant to see the new attribute. To provide fully assisted social IdM, suggestions for attribute allocations should be presented as well.

The algorithm is capable of classifying contacts into semantically associated groups based on their mutual contacts. However, while this classification is valid and corresponds to the user’s classification, as demonstrated by similarity values up to 100%, the user might still prefer to share dif-

ferent information with members of one group. For example, there may be both loose acquaintances met in a student club, as well as members that have grown to be deep friends. This is illustrated in Figure 6, in which the circles represent classes of attribute values, with the most intimate values, such as sexual preferences, located in the center and the least intimate values located in the outer circle. Contacts, represented by dots are placed on those circles corresponding to the user's willingness to share personal information with them. As one can see, groups like group A cause problems for audience segregation, as they cover multiple attribute classes. Groups similar to group B are more suitable, as they only cover a small range of attribute classes.

6. CONCLUSIONS

The pervasiveness of the social web poses many challenges for future research, especially in the field of privacy. Unlike the real-world, where personal information is ephemeral, in the online-world, this information is almost infinitely available while new information is added constantly increasing the existing complexity of managing different identities consistently. This permanency of information poses a great challenge for personal social IdM, since we are no longer free in constructing our identities because contradictory information may be available online. While the social web provides the platform for social IdM to everyone, regardless their technical expertise, privacy controls to raise awareness of implications using the social web and support the users in constructing and managing different identities are still in their infancy.

In this paper we propose assisted social IdM as a means to advance from the current state. In more detail, we provide a conceptualization of current and future social IdM and its formal foundations. Building on that, we propose several approaches to segregate audiences, which is a necessary step to disclose personal information within its belonging context. Our proposal is backed by a prototypical implementation and an evaluation that prove the applicability of the approach.

Acknowledgements

The research leading to these results is partly funded by the European Union within the Padgets project under grant agreement no. 248920.

7. REFERENCES

- [1] M. Bauer, M. Meints, and M. Hansen. Structured overview on prototypes and concepts of identity management systems. http://www.fidis.net/fileadmin/fidis/deliverables/fidis-wp3-del3.1.overview_on_IMS.final.pdf, September 2005.
- [2] H. Buhl and G. Müller. Der "gläserne Bürger" im Web 2.0 - Herausforderungen des "virtuellen Striptease". *Wirtschaftsinformatik*, 4:193–197, 2010.
- [3] K. Cameron. The laws of identity. <http://www.identityblog.com/stories/2005/05/13/TheLawsOfIdentity.pdf>, March 2005.
- [4] S. Claubeta, D. Kesdogan, and T. Kölsch. Privacy enhancing identity management: protection against re-identification and profiling. In *Proc. of the 2005 Workshop on Digital Identity Management*, 2005.
- [5] European Parliament. EU-Directive 95/46/EC. *Official Journal of the European Communities*, 1995.
- [6] E. Goffman. *The Presentation of Self in Everyday Life*. Anchor, 1959.
- [7] S. Gürses and B. Berendt. The Social Web and Privacy: Practice, Reciprocity and Conflicts in Social Networks. In *Privacy-Aware Knowledge Discovery: Novel Applications and New Techniques*. Ferrari, Elena and Bonchi, Francesco, 2009.
- [8] M. Hansen and M. Meints. Digitale Identitäten - Überblick und aktuelle Trends. *Datenschutz und Datensicherheit - DuD*, 30:543–547, 2006.
- [9] P. Jaccard. Étude comparative de la distribution florale dans une portion des alpes et des jura. *Bulletin del la Société Vaudoise des Sciences Naturelles*, 37:547–579, 1901.
- [10] R. Leenes. Context Is Everything Sociality and Privacy in Online Social Network Sites. In *Privacy and Identity Management for Life*. Springer, 2010.
- [11] B. Malin. k-Unlinkability: A privacy protection model for distributed data. *Data & Knowledge Engineering*, 64:294–311, 2008.
- [12] H. F. Nissenbaum. Privacy as Contextual Integrity. *Washington Law Review*, 79:119–157, 2004.
- [13] I. Oomen and R. Leenes. Privacy risk perceptions and privacy protection strategies. In *Policies and Research in Identity Management*. Springer, 2008.
- [14] A. Pfitzmann and M. Hansen. A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management, 2010.
- [15] PrimeLife. D1.2.1 - Privacy Enabled Communities. http://www.primelife.eu/images/stories/deliverables/d1.2.1-10.04.23-privacy_enabled_communities-public.pdf, 2010.
- [16] C. Shirky. Here comes everybody: The power of organizing without organizations. 2008.
- [17] Z. Tufekci. Can You See Me Now? Audience and Disclosure Regulation in Online Social Network Sites. *Bulletin of Science, Technology & Society*, 11:544–564, 2008.
- [18] S. D. Warren and L. D. Brandeis. The right to privacy. *Harvard Law Review*, 4:193–220, 1890.
- [19] E. Weiser. *The Functions of Internet Use and Their Social, Psychological, and Interpersonal Consequences*. PhD thesis, Texas Tech University, 2001.
- [20] A. Westin. *Privacy and freedom*. The Bodley Head, 1967.
- [21] I. H. Witten and E. Frank. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, 2005.