

Norm-based Access Control

Onuralp Ulusoy
Utrecht University
Utrecht, the Netherlands
o.ulusoy@uu.nl

Pınar Yolum
Utrecht University
Utrecht, the Netherlands
p.yolum@uu.nl

ABSTRACT

Collaborative systems, such as online social networks or Internet of Things, host vast amounts of content that is created and manipulated by multiple users. Co-edited documents or group pictures are prime examples of such *co-owned* content. Respecting privacy of users in collaborative systems is difficult because the co-owners of the shared content can have conflicting access policies about the content. To address this problem, recent approaches employ group decision making techniques, such as auctions. With these approaches, when a content is to be shared, all co-owners express their privacy preferences through the mechanism (e.g., by bidding) and the group decision mechanism reaches a decision to enable or deny access to the content. However, such mechanisms have to be carried out per content, making them impractical for most realistic settings. We argue that rather than employing a group decision mechanism on each content separately, it is more practical to watch for *privacy norms* that emerge in systems and make decisions using these norms, when possible. This paper borrows ideas from philosophy to represent privacy norms and develops algorithms to compute them in collaborative systems. We show that when privacy norms are identified correctly, they can enable collaborative systems respect users' privacy as well as decrease the need to engage in a group decision mechanism considerably.

CCS CONCEPTS

• Security and privacy → Privacy protections.

KEYWORDS

Privacy, Multiagent systems, Norms

ACM Reference Format:

Onuralp Ulusoy and Pınar Yolum. 2020. Norm-based Access Control. In *25th ACM Symposium on Access Control Models and Technologies (SACMAT'20)*, June 10–12, 2020, Barcelona, Spain. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3381991.3395601>

1 INTRODUCTION

Many recent software systems are built on the idea of collaborative computing, where multiple users share, manipulate and manage information about themselves as well as others. Online social networks (OSNs) are a prime example, where a user can put up a group picture

without explicit consent from individuals in the picture, and others can access to, comment on or even reshare the content, making it more visible to the world. For many users, this means that their personal lives are easily accessible to individuals or companies without them knowing about it.

Even though we are living in a privacy-conscious era with various policies in place to attempt preserving privacy, existing techniques have not been sufficient to detect, let alone handle these privacy violations. The main reason behind this is that privacy has been simplified to an informed consent, where the main assumption is that a user is in control of her data and chooses how to manage her privacy by giving appropriate consent. General Data Privacy Regulation (GDPR) [1] is an important policy, which is based on this idea of informed consent. Put simply, a user of a website is notified what kind of information will be collected about the user, for what purposes, and whether the information will further be processed. The user then gives an informed consent as to how her personal data will be shared. As a result, with appropriate techniques, it is possible to detect whether personal data have been leaked without the person's consent or used against her will. While GDPR assumes that each user can independently manage the privacy of personal data, the content that exists on collaborative systems, such as co-edited documents or group pictures, do not always belong to a single person. Further, many times content about a user is shared by others, not by the person herself. For example, a co-author of a jointly edited document can send a link of the document to whomever she sees fit. Or a user on an online social network can share a group picture publicly without explicit consent from those in the picture. It is possible that the individuals that are related to the content might have different and possibly contradictory privacy preferences [15, 31]. In these situations, when the sharing party is assumed to own the content, only her privacy preferences will be in effect. However, in many situations, the content might be *co-owned* by others that bring about the content in the first place (e.g., co-editors of a document). Hence, it is not sufficient to allow access to the content by only considering the sharing party's privacy preferences.

GDPR does not address how to tackle the privacy of content that pertains to more than one individual or that is shared by others about the user [5, 23]. We need to think of privacy for co-owned content different than the privacy of personal data since the sharing intentions and privacy preferences of all involved are at stake. Various collaborative privacy management mechanisms exist to tackle this. Negotiation-based agreement techniques for access decisions in OSNs [16, 31], propose that individuals negotiate themselves or through their software on how to share a given content. Argumentation-based techniques enable users' agents to exchange arguments to convince the others to comply with their privacy requirements [17]). Auction-based techniques let users reflect their privacy preferences on a content through bids and make decisions

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SACMAT '20, June 10–12, 2020, Barcelona, Spain

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7568-9/20/06...\$15.00

<https://doi.org/10.1145/3381991.3395601>

based on their outcome [29, 35]. These approaches are promising when each user is fully aware of her privacy expectations and can actively participate in the decision making whenever an access decision needs to be made. However, this is unrealistic for many systems where huge amounts of co-owned data are shared frequently but many users do not engage in configuring their privacy settings. Thus, it would be useful to be able to configure the access settings of a content without explicitly involving all the co-owners of the content into decision making.

Human societies often use norms for decision making [14]. We advocate norms as the basis of access control for collaborative privacy decisions. If the systems can identify the existing norms, then decisions can be made using them. This implies that a more complex decision mechanism, such as an auction or a negotiation, would not be required, speeding up the decisions that can be taken considerably. But, norms can also serve another important function. When an individual does not have or cannot formulate her privacy preferences, then the norms of a society can shed a light as to what is appropriate. The user does not have to follow the social norms at all time. If the user does not want to follow the norms or none of the existing norms apply to a given situation, then the system can still use a collaborative privacy management mechanism to make a decision. Contrary to successful access control schemes, such as role-based access control [25] or relation-based access control [12] that mostly apply on content that is owned by a single individual, norm-based access control enables access decisions on co-owned content.

This paper describes the principles of norm-based access control and develops an approach named PRINOR where access control decisions can be taken based on the *norms* that are generated from the previous privacy decisions in the system. We represent the different privacy norms in OSNs using Tuomela’s categorization [36] and develop algorithms to identify these norms in a given system. The usage of the algorithms enables users to choose between enforcing personal privacy settings and following the norms in the system. We show over multiagent simulations that when privacy norms emerge, they can be used in place of collaborative decision mechanism that require interactions for each content. Our analysis shows that the variations in the privacy expectations of the users have little effect on the success of PRINOR. We also apply PRINOR on a case study with real-life social network and image content data sets to demonstrate norm emergence and privacy decisions.

2 PRIVACY NORMS

We study the representation, emergence and usage of norms in collaborative systems, where a set of users are related to each other through a set of relations types (r_{type}), such as friend, colleague, and so on [12]. Each user can share content that pertains to herself as well as others. A user’s privacy preference about a content could depend on the properties of the content as well as the relation types with whom the content is shared. For example, a user might not want her holiday pictures to be shown to colleagues, but might be fine with work pictures to be shown. When the co-owners have conflicting privacy preferences, they need to reach a *privacy decision* that states if and how the content will be shared.

Each user in PRINOR is represented by a software agent, which keeps track of the privacy expectation of its user for sharing content [17]. We represent contents with a content descriptor c_{des} , which is a set of two tuples (x, n) , where x is a context such as holiday, work, and so on and n is the percentage of how much this content belongs to x . For example, a picture taken at a bar might be represented as: $\{(nightlife, 77\%), (leisure, 12\%)\}$. The context information might be available in the system but it could also be derived automatically as we explain in Section 6 through software that produces tags and confidence intervals. Depending on the content, the set might have more tuples. We do not require the sum of the percentages to be equal to 100% since the content may highly be relevant to multiple contexts making their sum way over 100%. Alternatively, we may not have enough evidence to associate a content with contexts; hence the sum may be less than 100%. For each content c , we also specify the set of co-owners c_{own} , whose privacy is possibly being affected by the content and thus should have a say about content’s privacy decision. In general, if the content is a picture, c_{own} could consist of users tagged in the picture or if the content is a co-edited document, it could be the co-authors. We assume that this set can be retrieved from the system as is the case with most collaborative systems.

PRINOR contains norms to capture the privacy preferences. Informally, privacy norms capture the common behavior for accessing a particular type of content with a particular set of users. In most domains, it is generally assumed that when the actions of the agents are in conflict, norms that are fully applicable to all the agents can be found. A prime example is the well-studied traffic domain [22], where norms such as driving on the right side of the road might emerge because it can be observed that mixed usage of sides leads to accidents. However, norms have to be rethought in the case of privacy. Privacy norms, by definition, are different from other norms because it is extremely difficult to find privacy norms that could satisfy the expectations of all of the agents. That is, there is no one right norm to make everyone happy. Whereas in a domain such as traffic, an agent benefits from obeying an established norm; in a privacy related domain, complying with a social norm might harm the privacy of agents, depending on their privacy requirements. Therefore, norms should be allowed to emerge at different levels (e.g., norms in a group of users, norms for a specific relation type, and so on) and the norms should be evaluated continuously to find out if they still fit to the expectations of the population. Moreover, an agent should always be allowed not to follow a norm so that privacy preferences of individuals and minority groups can still be respected.

Tuomela [34] categorizes norms as social and personal norms. Social norms are formed according to the behavior of the society, and can be sanctioned if one does not comply with them. Personal norms are related to individuals’ comprehension of the environment, and their beliefs about which actions are right or wrong within the society. Tuomela further divides social norms to r-norms (rule norms) and s-norms (social norms), and personal norms to m-norms (moral norms) and p-norms (prudential norms). We adopt this classification to model privacy expectations as norms and formally represent it similar to existing formalisms [3, 22, 36], such that a set of pre-conditions determine the activation of a sharing action to be taken. We also aim to handle context-based privacy preferences [4] with our norm definitions. Since our focus is more on the emergence of

norms rather than their violation, we do not include norm sanctions explicitly [26]. We consider social norms to be governed by an overseer mechanism (e.g., an OSN provider), while personal norms are handled in a distributed manner.

r-norms are imposed by an authority to the individuals. These are simply laws of a collaborative system, without leaving any room for personal choices, e.g., OSN denies access to any violent content. In PRINOR, an *r-norm* is a 2-tuple structure represented as $r_i = \langle c_{des}, act\{deny\} \rangle$, where c_{des} is the descriptor of the content on which this *r-norm* applies and action is the sharing decision, which in this case *deny* for the related descriptor.

s-norms are related to the common understanding of the society that apply to every individual. For example, in a given society, a norm of not sharing content that contains alcohol might emerge. *s-norms* are 3-tuple norms represented as $s_i = \langle r_{type}, c_{des}, act\{access, deny\} \rangle$, where r_{type} is the relationship type between the co-owners for a content, c_{des} is the descriptor for which the *s-norm* will apply and *act* is the assigned action of the norm, which could be either enabling or denying access to the content. *s-norms* are emergent norms depending on the previous collaborative decisions within the social network. We employ r_{type} since *s-norms* are generated according to an overview of the societal decisions, we aim to conceal the specific actions of individuals to the mechanism that can generate *s-norms*, and only reveal generic relationship types that the privacy decisions apply to.

m-norms capture an individual's own privacy preferences. These are moral norms that individual agents store for their future privacy related decisions. The representation of *m-norms* are identical to that of *s-norm* ($m_i = \langle r_{type}, c_{des}, act\{access, deny\} \rangle$), though *s-norms* emerge over time and calculated by PRINOR, whereas *m-norms* are given norms of an agent. *S-norms* exist as part of the collaborative system, whereas *m-norms* are private to each agent.

p-norms are defined as what individuals understand as the rational actions. For example, a group of agents might always share their co-edited documents with others. In this regard, prudential norms are useful for exploring normative behavior within specific sets of agents. A *p-norm* is a 3-tuple $p_i = \langle c_{own}, c_{des}, act\{access, deny\} \rangle$, where c_{own} is the owners of the content that is described with c_{des} and the action is to enable or deny access to content.

3 PRINOR

The above norm types are all stored in respective norm bases. Initially, each agent has a personal *m-norm* base, which can only be updated by the agent itself. An *m-norms* base can be thought as the privacy policy of the agent. For now, we assume that the *m-norm* base for an agent does not change over time. At the beginning, the collaborative system itself has a single *r-norm* base that contains all the laws of the system. The norms in the *r-norm* base are stored and updated by the system provider itself. Again, we assume that the *r-norms* are static and do not change over time.

S-norm base contains the social norms in the system. These social norms emerge based on the privacy decisions made in the system by the individual agents. That is, the agents themselves change the understanding of privacy in the system and contribute to formation of norms. There is also a single *s-norm* base in the system but it

is updated over time. Sometimes privacy norms can emerge at the society level, but sometimes at a smaller, group level as a *p-norm*. A group can be two or more agents that have shared a content at one point in time. A *p-norm* base stores unique group related norms, therefore each agent stores their own *p-norm* base for the groups she has been in, and updates it according to the given specific group's previous privacy decisions that were made with the employed collaborative privacy mechanism. Contrary to *s-norm* base, *p-norm* base is distributed. Because of this it is possible that some agents in a group may not reach an emergent *p-norm* for an upcoming decision due to the differences in the subsets. We resolve this by enabling one agent to inform all the others in the agent set when a new *p-norm* emerges, and others update their *p-norm* base accordingly.

PRINOR works as follows: When an agent wants to share a content, which is co-owned by other agents, the uploader agent checks if it is desirable for all the co-owners to share the content, considering the norms. This is done by considering the type of the content and the relationship with other co-owners. Since four types of norms are in effect, there can easily be conflicts among various norm-bases. For example, an agent's *m-norm* might permit sharing a content publicly, whereas the *s-norm* in the system might prescribe otherwise. This calls for an ordering of norm bases. Dechesne *et al.* [11] show that there are several individual characteristics that affect the decision process of the individuals, such as compliance with the law, abiding to social conventions or behaving according to individual preferences. An individualistic agent might first check its *m-norm* base and refer to other bases only if there are no related norms in this base. A social agent can prefer to put *s-norms* in front of *m-norms* while a law abiding agent always places *r-norms* at the top. An interesting choice question comes up with *p-norms* and *s-norms*, since they both are in the social context, while the former only includes a specific set of agents that the agent directly has a relationship with. In this work, we assume that *p-norms* always dominate *s-norms*, since norms within direct relationships represent more precise behavior than the norms emerging from a community which is formed by indirect relationships (i.e., agents that do not have a relationship, but are present in the same OSN community) and that *s-norms* dominate the *m-norms* since we are interested in understanding the benefits of making privacy decisions using societal norms. Using this ordering, the uploader agent checks its *r-norm*, *p-norm* and *s-norm* bases to see if a norm matching with the content descriptor exists. If so, it is applied. It might be the case that none of the norms in the norm bases are applicable. If so, the agents engage in a decision mechanism, such as auctions or negotiation, and the final decision is made according to the chosen collaborative privacy mechanism. When agents engage in a decision mechanism, they use their *m-norm* bases to reveal their valuations. If such a mechanism is used, then the outcome of the mechanism also updates the *p-norm* base of the co-owner agents and *s-norm* base of the OSN, where new possible norms can be formed for future incoming co-owned content.

Figure 1 depicts how PRINOR works when a decision is to be made for an incoming content for three agents, Alice, Bob and Carol, who have a friendship relation. The legend on the left side describes the norms and the contents relevant in the system. The numbers of arrows indicate the order of actions. Alice wants to share the contents but the contents are co-owned by all three agents. We give various examples of how PRINOR would work on an incoming content. In

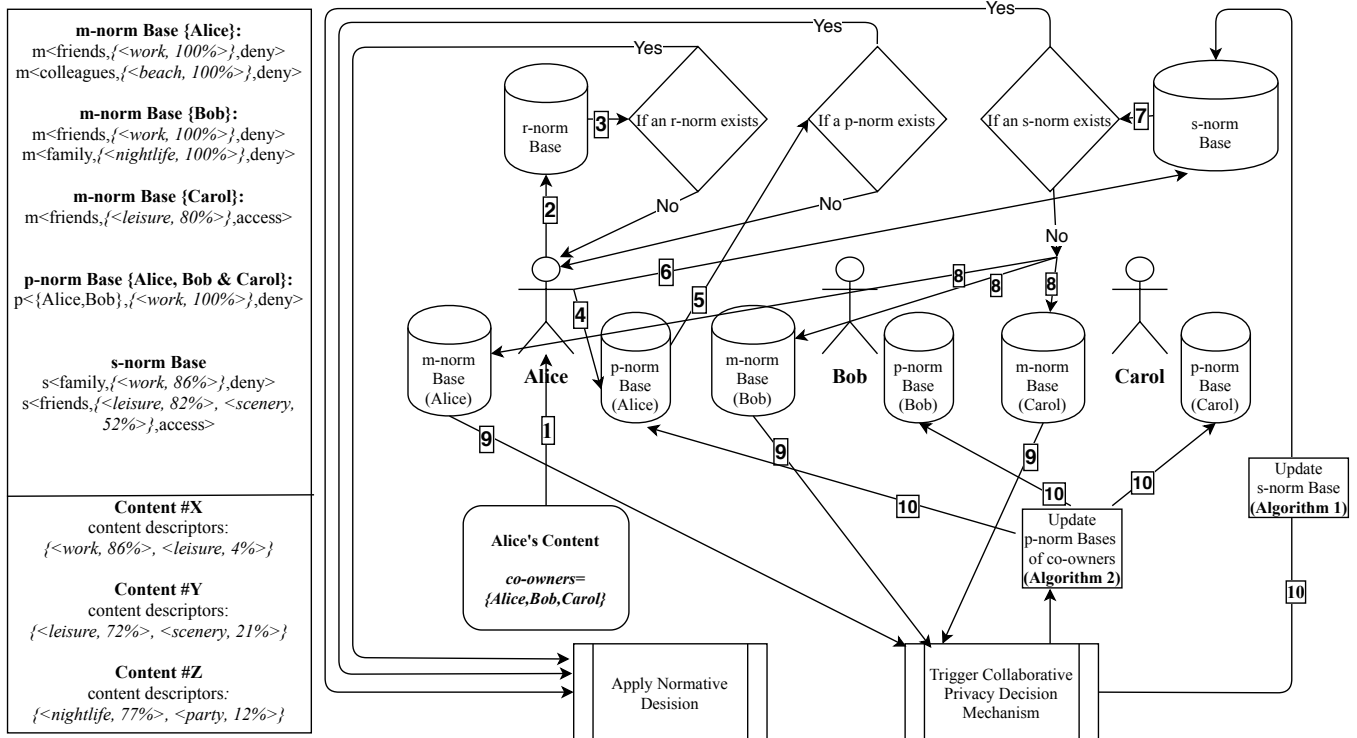


Figure 1: Normative decision mechanism process for an incoming content co-owned by three agents.

our examples, we exclude r-norm base checking for brevity, since it is only applied simply when a content type is forbidden by the OSN provider and always checked initially by all agents to not receive possible sanctions.

EXAMPLE 1. The incoming content X, which is mostly related to work context, according to the content descriptors. As the uploader, Alice checks her *p-norm* base, where all previous normative privacy mechanism decisions by every subset of these three agents are stored. *p-norm* base includes a fitting *p-norm* established between Alice and Bob, with *deny* action. Since this norm would be in effect in the greater group as well, the agents do not share the incoming content.

EXAMPLE 2. The incoming content is Y, and mostly related to leisure context. Since Alice does not have a related *p-norm* in her *p-norm* base, she checks the *s-norm* base and finds a similar *s-norm*, where the content descriptor indicates 82% relatedness with leisure context. Alice, Bob and Carol can comply with this *s-norm* to share the content according to *s-norm* base, since they are indeed in friendship relation.

EXAMPLE 3. For incoming content Z, Alice does not have an established *p-norm* in her *p-norm* base. Content Z does not fit into *s-norms* in the *s-norm* base, either. Therefore, collaborative privacy decision mechanism should be triggered, and the decision should be made according to agents' *m-norms*. Since this is a mechanism based decision, *p-norm* base of Alice, Bob and Carol is updated with the current decision. The *s-norm* base of the OSN registers this decision to be used for the generation of norms.

4 GENERATING NORMS

While a system starts with users' *m-norms* and the system's *r-norms*; *s-norms* and *p-norms* emerge over time based on the interactions of users. Further, an *s-norm* that emerges in a system may totally contradict the values of an agent as represented by an *m-norm*. Contrary to other domains where norms are to be followed by all, here for the privacy domain, we would like to give agents the option not to follow an *s-norm*. This necessitates a decision to follow or ignore an *s-norm*.

4.1 Identifying S-norms

Recall that each co-owned content in the system requires a privacy decision according to their contextual properties, and the outcome is to enable or deny access to the content with a set of relationship types. Given a set of such decisions, Algorithm 1 clusters the decisions to identify potential *s-norms*. Essentially, the algorithm places all the content over a multidimensional space according to their descriptor and the relationship type of the co-owner agents. This space contains all the decisions considering its various properties as dimensions. Then, we cluster this space such that each cluster contains content that have similar attributes. Finally, the clusters can be checked for being a possible *s-norm*. Since the evolution of social norms depend on many factors, continuous update of *s-norms* is essential to capture the current state of social normative behavior in the environment [26]. Therefore, the algorithm is run periodically in order to find out about new emerging norms or exclude norms that became obsolete over time.

OSNs enable users to continuously share tremendous amount of content. In a real life application, clustering every content in short periods would be infeasible, since it would require massive computing power. Therefore, a simple clustering algorithm which is sufficient enough to distinguish between contextual properties of content, in our case, the dimensions of the descriptors, is essential. In light of this intuition, we employ *k-means* algorithm to cluster content and then check all the clusters for normative behavior. *k-means* is a clustering method where n number of elements in a unidimensional or multidimensional space are partitioned into k clusters, where each element is assigned to the nearest mean of the elements in a cluster [32]. Determining the number of clusters is difficult. Rather than having a fixed number, we adjust it as needed. More precisely, in Algorithm 1, we start with a small number of clusters containing large amounts of content, hence we define a small k value resulting in a big n value. As a heuristic to determine normative behavior within a cluster, we use qualified majority to check the privacy decisions for the content within. According to qualified majority heuristic, we consider a cluster a candidate for being normative, if at least 66% of the privacy decisions are the same for the content in the cluster. If a candidate normative cluster is found with the initial k , it is saved to the *s-norm* base and the content within is removed from further calculations. For the remaining clusters that does not show normative behavior, k is increased and new clusters are formed to check if normative behavior emerges with smaller number of content within clusters. This approach continues until a threshold for the minimum number of agents in a cluster is reached, and the algorithm stops at that point to save the rest of the clusters as *non-normative*.

When applying naive *k-means* clustering, each content can be placed into the closest cluster, because all content is assumed to have the same dimensions. However, shared content in real life would have differing contextual properties. Since we take each contextual property as an additional dimension, the number of dimensions might become high. Moreover, many content would not have a common contextual property. A simple approach would be to still consider all possible contextual properties as separate dimensions, and assign the value of zero if a content is not related to this content descriptor. This would make the space rather sparse. With a large amount of dimensions, this can easily become infeasible since each content would have many dimensions valued at zero. As a result, the clustering can yield very distant clusters each containing only a small number of content. To resolve this, we propose a dimension reduction for the domains with a large variety of content descriptors. With this reduction, for a privacy decision of an incoming content, only the previous content that share descriptors are taken into consideration for clustering. Thus, the only dimensions required for clustering would become the incoming content's descriptors, which would significantly reduce the computation required for finding *s-norms*.

Another aspect to consider for social norms is the changes in the behavior of the society over time. As the time passes, the values of the people change, which also might cause some norms to become obsolete while new ones emerge. Thus, we employ an aging curve [37] for privacy actions, denoted as $R = e^{-t/S}$. Here, R is the *retrievability* of the privacy action, while t is time passed since the decision was taken and S is the *stability* of memory. According to

this equation, a recent privacy decision would take a value closer to 1, while over time it's value would be close to 0, and the speed of aging is dependent on the S value. Let us consider two examples of privacy decisions, one taken an hour ago for sharing a content while one taken 1000 hours ago for not sharing. If we define the S value according to a calendar month, hence, 720 hours, the first one would have a value of ~ 1 for retrievability, while the second is ~ 0.25 . If these were the only privacy decisions in the system, the naive qualified majority calculation for *s-norms* without considering the aging curve would have given 50% for sharing and 50% for not sharing a similar content. With aging curve in place we give the calculated R values as weights to the privacy decisions. Hence, the two example privacy decisions would result in 80% for sharing and 20% for not sharing, since the recent privacy decision is considered more significant for capturing current social behavior.

Algorithm 1: Generation of *s-norms*

Input: mk , minimum number of clusters
Input: t , threshold for min. number of agents in a cluster
Input: $pDec$, previous privacy decisions within OSN
Input: S , stability parameter for aging of decisions
Output: $cList$, a set of clusters generated from $pDec$

```

1 foreach  $item$  in  $pDec$  do
2    $R_{pDec} = \text{aging}(pDec, S)$ 
3 while  $pDec$  not empty do
4    $tempcList = \text{k-Means}(mk, pDec, R_{pDec})$ 
5   foreach  $cluster$  in  $tempcList$  do
6      $hasQualifiedMajority = \text{checkpDec}(cluster)$ 
7     if ( $hasQualifiedMajority = \text{true or}$ 
8        $size(cluster) < t$ ) then
9        $\text{add}(cluster, cList)$ 
10      foreach  $item$  in  $cluster$  do
11         $\text{remove}(item, pDec)$ 
12    $mk += 1$ 
13 return  $cList$ 

```

For each periodic call of *s-norm* base update, the minimum cluster count parameter (mk), the minimum size threshold parameter for a single cluster (t), all the previous privacy mechanism based decisions ($pDec$) and stability value (S) are taken as input for Algorithm 1. The algorithm starts with calculating retrievability values of all previous privacy decisions in *s-norm* base according to the stability parameter (lines 1 and 2). Then, for each item in $pDec$, a temporary list of clusters are assigned with *k-means* algorithm, where all items in $pDec$ are clustered into mk clusters. In line 5, a for loop begins, which checks the temporary cluster assignments, and determines if the cluster shows a normative behavior (i.e, qualified majority of the privacy decisions are the same), or the size of the cluster is below t value. If one of these conditions is satisfied for a temporary cluster, the cluster is added to $cList$ in line 9 and all the items of the cluster are removed from $pDec$, ending the iteration. If there are still remaining items in $pDec$, another iteration starts to determine new clusters, until all items from the initial $pDec$ are assigned to a cluster in $cList$ output.

4.2 Deciding to Follow an S-Norm

After *s-norms* are identified with Algorithm 1, they are stored in the *s-norm* base. Whenever an agent is making a privacy decision, it will check the *s-norm* base to see if any of the *s-norms* are applicable. If so, then the agent needs to decide if it would want to follow it.

To determine if a prescribed social norm of a cluster should be used as a privacy action for an upcoming content, agents can check three types of metrics. First, the percentage of the suggested normative privacy action for all privacy decisions in the closest cluster should still be in consideration, since a higher percentage would suggest homogeneous behavior of the society while a lower percentage indicate more heterogeneous behavior. Second, the distance of the content in consideration to the center of the cluster should be measured to understand how much the content is similar to the content present in the cluster. A content that can be placed closer to the center would mean that contextually, it is strongly correlated with the others. A third metric could be to check if the agent has established *m-norms* for similar type of content. If such norms exist and the privacy action is the same with the prescribed *s-norm* action, it would strengthen the agent's belief to comply with the *s-norm* while a different action would affect it negatively. We call these three types of metrics *majority percentage* (MP), *contextual similarity* (CS) and *decision similarity* (DS), respectively. All three metrics are defined to be between 0 and 1, and we apply an α weight in relation to these three metrics, again between 0 and 1, to compute a likelihood value to comply with the prescribed *s-norm* decision, which we abbreviate to *SD*.

The *majority percentage* for the *s-norm* privacy action is provided by Algorithm 1, which requires no further computation. To compute *contextual similarity*, we place the incoming content in the cluster and compute the Euclidean distance of every content descriptor dimension to the center. Then we do the same for the content in the same cluster that is furthest from the center. With the second distance, we normalize the first distance in a way that the furthest content would give *contextual similarity* value as 0 and the center itself would be 1. For example, in a single dimensional context, if the distance of the furthest content to the center is computed as 4, and the distance of the incoming content to the center is computed as 1, *contextual similarity* would be $((4 - 0) - 1)/4 = 0.75$, which means the content is strongly related to the cluster contextually. To compute the *decision similarity* metric, agents check if they have any *m-norms* for a similar type of content. If they do, then the *decision similarity* is simply the number of these moral norms divided to the number of all *m-norms* stored by the agent. This affects the decision for considering the incoming content normative positively, if the privacy decision of the considered *m-norms* are the same with the *s-norm*'s majority decision. If not, the effect of the *decision similarity* becomes negative.

With the weighted averages in consideration, the final decision to comply with the *s-norm* for the incoming content is shown below.

$$SD = \begin{cases} \alpha * (CS * MP) + (1 - \alpha) * DS, & \text{if } >0. \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Equation 1 divides the defined metrics into two parts according their relation with each other. Contextual similarity of the incoming content for the cluster, and the majority privacy decision of the

cluster are closely related to each other, hence we multiply them with each other, and also multiply this with α for weighting these parameters' importance. Decision similarity is related to differences of individual choices against the society, therefore it is weighted with $1 - \alpha$. With an α value closer to 1, the *SD* is more reliant on the society choices, while a value closer to 0 would give more importance to the similarity of individual privacy requirements with social norms. Referring back to the examples in Figure 1 of Section 3, we give an example below to depict how agents can use the *SD* metric to decide complying or denying social norms.

EXAMPLE 4. Consider Content Y from Figure 1 as the incoming content uploaded by Alice, and *MP* of the related cluster is 84% for not sharing action, given by Algorithm 1. For the calculation of *CS*, the content descriptor has two dimensions. Let us assume the furthest content of the same *s-norm* cluster has a content descriptor as <leisure, 32%>, <scenery, 13%> and each context type has the same importance. For the leisure contextual dimension, the furthest content of the cluster has 50% (82%-32%) and 39% (52%-13%) distance for the scenery context. The same values for Content Y are 10% (82%-72%) and 31% (52%-21%), respectively. Therefore, *CS* value is the mean of $(50\% - 10\%) / 50\%$ and $(39\% - 31\%) / 39\%$, equalling to ~50%. Alice does not have any *m-norms* related to the *s-norm* in consideration in her *m-norm* base, which is consisting of two *m-norms*. Thus, the *DS* value will be 0 (0/2). If the α is given as 0.8, *SD*, which is the likelihood of complying with the norm would be calculated as ~33% $(0.8 * (0.84 * 0.5) + 0.2 * 0)$, which would prescribe Alice to follow the norms with a one in three probability.

4.3 Identifying p-Norms

Prudential norms are the second type of societal norms in our mechanism. The *p-norms* only bind the users in the group and not the society as a whole. Hence, we require that *p-norms* are kept and updated separately by each agent (rather than by the OSN provider as was the case with *s-norms*). Essentially, *p-norms* represent previous collaborative privacy mechanism based decisions of co-owner agents for a content. To keep privacy requirements simple, agents only classify *p-norms* according to the major content types (i.e content type with the highest relatedness value). In addition, agents keep track of the co-owner IDs, since *p-norms* are the norms that emerge between specific sets of co-owners. Algorithm 2 shows how an agent generates a *p-norm* after deciding on an incoming content.

After an incoming content where a privacy decision is required, Algorithm 2 is triggered by each co-owner agent of the content to generate a *p-norm*. The inputs include the major content type (*mct*) of the content, which defines the highest valued content type. *aDec* contains previous privacy decisions of co-owners, including the decisions made by subsets of the co-owners. This enables the algorithm to propagate previous privacy decisions of smaller subsets of co-owners to the entire set of co-owners. Since a subset of co-owners might not fully represent the behavior of a bigger co-owner group, we introduce a difference parameter (*d*) in the algorithm, which enables the system to adjust the impact of previous privacy decisions with different size of subsets of co-owners. The algorithm starts with assigning counts of each action type possible for a privacy decision as zero (line 1). Then for each item in the *p-norm* base of an agent, the algorithm counts the previous privacy actions, where

Algorithm 2: Generation of p -norms

Input: c , content in discussion
Input: co , list of co-owner agents for c
Input: $c(mct)$, major content type of c
Input: d , difference parameter for co-owner similarity
Input: $aDec$, agent's previous privacy decision list
Input: qMP , qualified majority percentage threshold
Output: $pList$, a list of p -norms, forming the p -norm base of the system

```
1 initialize actionType counts as zero
2 foreach item in aDec do
3   if ( $\forall$  item(co-owner) in co) then
4     dif = (size(co) - size(item(co-owners)))
5     foreach act in actionType do
6       if (item(privAction) eq act and c(mct) eq item(mct))
7         then
8           count(act) +=  $1/d^{dif}$ 
9           totalCount +=  $1/d^{dif}$ 
9 foreach act in actionType do
10  if (count(act)/totalCount > qMP) then
11  c(privAction) = act
12 update_pList(p<co,c(mct),c(privAction)>)
13 return pList
```

the major content type is the same as the current content, and all the co-owners of the item in p -norm base are elements of the co-owners set of the content. In line 4, the difference between the size of co-owners of p -norm base item and the size of co-owners for the content in consideration is computed. For example, if the content has three co-owners named Alice, Bob and Carol; and the p -norm base item has Alice and Bob as co-owners, then the distance is simply computed as $3-2 = 1$. Then the count of each action is increased according to the formula on line 7. With the same example above, if the difference parameter d was assigned 2, the increase would be computed as $1/2^1 = 0.5$, reducing the effect of it from a p -norm base item which has all three of the co-owners of the content in consideration. After all action type counts have been computed, another loop checks the action types to decide if a normative behavior exists. This comparison is made according to qualified majority percentage threshold (i.e., 66%), which can be set as input (qMP). If an action type percentage is above the threshold, agents consider this as normative behavior. Notice that co-owner subsets might be different for groups (e.g., for a content co-owned by Alice, Bob and Carol, previous decisions established between Alice and Bob are not known by Carol), yielding some agents not be aware of an existing p -norm. We remedy this by requiring each agent to notify others of p -norm updates. In order to synchronize the p -norm bases between co-owners, every agent informs the other co-owners when a new norm emerges and the others update their own base if they have not already reached the same norm. Finally, p -norm bases of all the co-owners are updated with the new p -norm. The agents can choose to apply the p -norm or make a decision with collaborative privacy mechanism using their m -norms.

5 EVALUATION

We evaluate PRINOR in a multiagent simulation environment that we developed in Java¹. Each agent in the simulation represents a user. The users, and thus the agents are related to each other through one relationship. Each agent is uniquely identified with an identifier. Each agent has a set of m -norms that are generated automatically. Each content in the OSN is assigned a descriptor. In real life, this information would come from the features or tags of the content. Here, we assume that the descriptor is available. For n number of content type categories, a content is placed in an n -dimensional space which enables the mechanism to both find out similar content types and match privacy requirements of agents with the content in consideration. In addition, a content has a set of co-owners, which are the agents that are within the OSN that have some of their private information represented in the content.

We include 100 agents and 10000 contents for each of our simulation runs, where each content is randomly assigned to 2 to 5 co-owners, and a descriptor with 4 elements, while each element is a two-tuple with a context and a value between 0 and 100, representing the significance of the content to the given type, 100 being the most. We represent each agent's privacy requirements with m -norms, while the simulation checks the evolution of p -norms and s -norms. We exclude r -norms from the simulations as our focus is on the correct emergence of p -norms and s -norms. On each simulation, one content is introduced to the mechanism sequentially. First, the societal norms are checked to reach a decision. If relevant societal norms are not present, then the decision is made according to the m -norms of the agents. For m -norm based decisions, our current mechanism allows us to employ different mechanisms such as auctions, negotiation or argumentation. However, these mechanisms require rather complex computation. In order to keep computational complexity low, we employ majority voting as the collaborative privacy mechanism in our evaluations. With this simulation setup, we answer the following questions: (i) Do s -norms and p -norms emerge over time and if so, what percentage of access control decisions are taken by these norms? (ii) Do the norms that emerge enable agents to make correct access control decisions?

5.1 Emergence of Social Norms

In a pioneering work on norm emergence, Sen and Airiau [27] show that norms emerge even when the population size and heterogeneity vary. Following this, we introduce a *homogeneity* variable to capture how much of a society has similar privacy understandings. In our approach, if the homogeneity of the society is 0, then all the agents in the population can have different privacy choices. We run our simulations for investigating emergent social norms with different levels of homogeneity. We achieve this by making a subset of agents having the same action type for a given type of content in their m -norm base, while the rest is assigned a random type of action for their m -norms. Our homogeneity levels are 0%, 10%, 20%, 30%, 40%, 50%, 75% and 100% respectively, 0% representing full random behavior and 100% full homogeneity. The reason of having bigger margin between the latter three levels is that societal decisions are almost similar when homogeneity percentage is bigger than 50% in the network. The simulation starts forming s -norms using Algorithm 1

¹Repository link for our simulation: <https://git.science.uu.nl/o.ulusoy/PriNorSim>

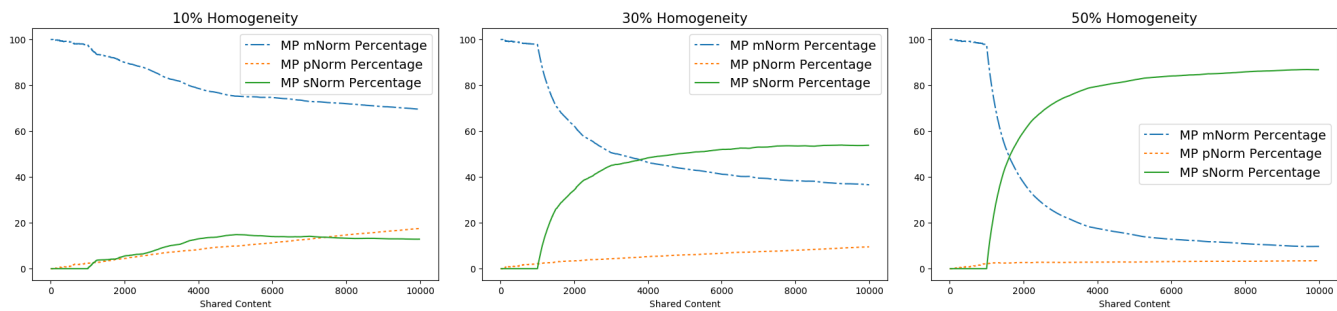


Figure 2: Percentage of norm types over different levels of homogeneity for qualified majority s-norm decision.

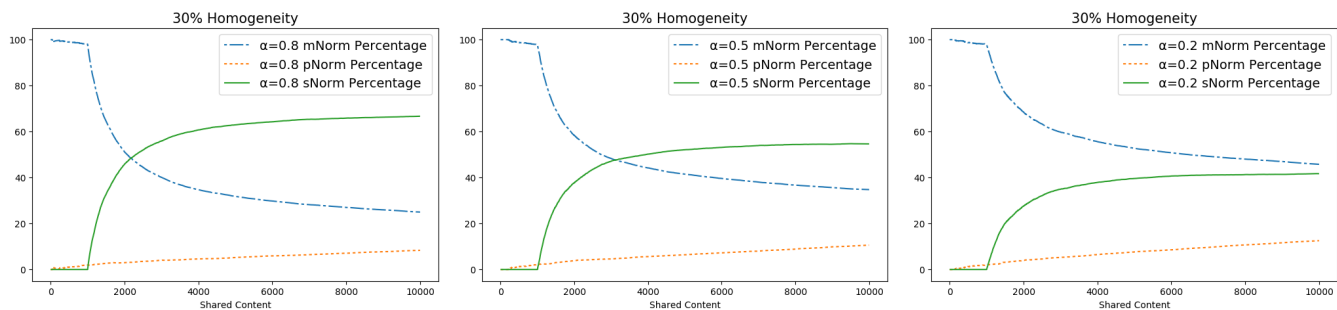


Figure 3: Percentage of norm types over 30% homogeneity based on likelihood of following the norms function for various α values.

after 1000th content shared in the OSN and reruns it after every 250 content for updating the *s-norm* base of the OSN. To compare how agents decide to follow the norms, we evaluate four different setups. In the first setup, agents follow *s-norms* at least when qualified majority percentage for a single privacy action is satisfied. The other three setups employ the *SD* formula, which gives a likelihood value of following the norms for agents, with three different α values, 0.8, 0.5 and 0.2 respectively. For each homogeneity level combined with each of the four setups, we run 5 simulations and measure the percentage of decisions taken with *m-norms*, *p-norms* and *s-norms*.

Figure 2 plots the percentage of decisions that are taken by *m-norms*, *p-norms* and *s-norms* as new content is introduced to the system for populations with 10%, 30% and 50% homogeneity and when the *s-norm* decisions are made according to qualified majority percentages of *s-norm* clusters. For 10% homogeneity, 12.88% of all decisions were made with *s-norms* while 17.54% of all 10000 content is decided according to *p-norms*, without the need of triggering the collaborative privacy decision mechanism with *m-norms*. This can be seen as a significant improvement, since our norm based method reduces the need to trigger a decision mechanism by ~ 30 percent, even when a tiny fraction of the society behaves homogeneous and the amount of co-owned content is sparse. The sparsity comes from having 100 agents randomly assigned as co-owners of 10000 content, since the same subset of agents can only have a very limited number of content with the same major relation type and content type. Therefore, building up *p-norm* base of every agent becomes

a difficult task with the limited previous knowledge about the co-owned content with the same subset of related agents. With 30% homogeneity, more than 63% of the decisions can be made with *p-norms* and *s-norms* and with 50% homogeneity, the necessity of applying a privacy decision algorithm with *m-norms* goes below 10%. Our results show that even if a small amount of agents in a system act similar instead of randomly behaving, social norms can emerge and effectively be used for collaborative access control decisions.

Even though deciding only according to qualified majority decisions for *s-norms* reduce the need of a collaborative decision mechanism significantly, the emergent norms might differ from the privacy understanding of individual agents. Some agents might act differently than the society, therefore applying social norms might create privacy decisions that the agents would not want to achieve by themselves. To account for this, we introduced a likelihood to follow *s-norms* formula (*SD*) in Section 4. We evaluate the *SD* formula with three α values and plot the results for 30% homogeneity in Figure 3. We omit the rest of the evaluations with different homogeneity for brevity, since all levels show similar behavior in comparison with qualified majority based decisions and this homogeneity reflects the real life social behavior more than the both sides of extreme levels.

Recall that when α is high, agents assign a high weight to a given *s-norm* but value their own *m-norm* about a content less, if such a norm exists. Accordingly, one would expect that with high values of α , more decisions would be taken with *s-norms* and with low values, the number of decisions would decrease. Our results confirm this. The results show that when α is set to 0.8, the decisions based

on *s-norms* takes up two thirds of all, and with *p-norms* the total norm based decisions constitute ~75% of the privacy decisions. The number of *s-norm* based decisions decrease with $\alpha = 0.5$ setup almost to the number in qualified majority percentage setup. The number of decisions is even fewer when α is assigned as 0.2, but still reduces the need of mechanism based decisions to less than 46%, when combined with *p-norms*. However, this decrease can still be beneficial for the agents, since they ensure that the applied social norms are in line with their own privacy understanding, while rejected norms are quite different than theirs. This brings up the *correctness* of the applied social norms into question, which we will investigate in the next subsection.

5.2 Correctness of Social Norms

Usage of norms decrease the need of a complex privacy decision mechanism, but do they lead to correct privacy decisions? We measure correctness by comparing norm-based decisions with collaborative privacy decision mechanism results. If the outcome of the norm based decision is the same with what the mechanism would give, we consider it as a correct decision. Since our current setup enables the simulation to evaluate both emergence and correctness of the norms within the same run, we investigate the correctness aspect of PRINOR with our multiple runs for various homogeneity and α values executed for Subsection 5.1 and present our findings about it in this subsection.

Table 1 shows the percentage of *s-norm* and *p-norm* decisions over all our setups with various homogeneity levels and agent decision types to follow *s-norms*, along with their correctness ratios. An immediate result is that in any setup, decisions made using *s-norms* are at least 75% correct (HL=0%, $\alpha=0.8$). The percentage of correct *s-norm* assignments increases with higher homogeneity levels, and end up at 100% when all agents in the community are homogeneous in their privacy actions. When we compare different *SD* setups, we observe that the highest correctness percentage comes with $\alpha = 0.2$ parameter. This is an expected outcome since with lower α values, the agents mostly follow the social norms when they are in line with their own privacy policies. $\alpha = 0.8$ setup with the *SD* metric performs the best with lower homogeneity levels to reach a high number of emergent norms, while keeping a reasonably high correctness ratio. Qualified majority setup has the highest *s-norm* percentages with the highest homogeneity levels, since almost all the agents behave the same.

Our results indicate that with a fine-tuned setup, even in unrealistically low homogeneity levels, ~90 percent of the entire *s-norm* based decisions are correct. For example, when α is 0.2 and the homogeneity level is %20, PRINOR can make %41.68 of the decisions for 10000 content with %91.42 correctness for *s-norms*. This means that ~4168 privacy decisions are taken with *s-norms* without any effort or feedback from the OSN users, and ~357 of the decisions were not correct, which is ~3.6 of the entire decisions. Note that emergent norms do not always make correct decisions. However, when the user does not know her privacy preferences or the number of decisions that need to be taken are large, they provide a suitable mechanism to make decisions. The choice of relying on the norms versus a complex decision mechanism can be decided by the user by setting the the α parameter, where a low value of α favors the user's own preferences

HL	SD	s-norm %	correct s-norm %	p-norm %	correct p-norm %
%0	MP	5.66	78.68	18.77	98.65
	$\alpha = 0.8$	60.98	74.95	9.69	98.76
	$\alpha = 0.5$	48.20	81.88	11.86	98.67
	$\alpha = 0.2$	36.73	87.34	13.84	98.72
%10	MP	12.88	78.81	17.54	98.94
	$\alpha = 0.8$	62.41	77.75	9.40	98.77
	$\alpha = 0.5$	50.19	83.77	11.28	98.67
	$\alpha = 0.2$	38.57	89.00	13.22	98.87
%20	MP	30.68	82.56	13.92	98.94
	$\alpha = 0.8$	64.66	81.85	9.05	98.67
	$\alpha = 0.5$	50.38	84.48	11.56	98.74
	$\alpha = 0.2$	41.48	90.52	12.93	98.49
%30	MP	53.86	82.81	9.53	98.96
	$\alpha = 0.8$	66.68	84.47	8.34	98.81
	$\alpha = 0.5$	54.64	89.38	10.60	98.54
	$\alpha = 0.2$	41.68	91.42	12.55	98.56
%40	MP	71.79	83.61	6.49	98.74
	$\alpha = 0.8$	68.43	86.49	8.09	98.74
	$\alpha = 0.5$	58.16	91.16	10.35	99.22
	$\alpha = 0.2$	47.07	94.47	11.78	98.51
%50	MP	86.84	87.41	3.46	99.26
	$\alpha = 0.8$	70.66	88.42	7.60	98.92
	$\alpha = 0.5$	60.00	92.45	9.60	99.07
	$\alpha = 0.2$	49.34	95.48	11.71	98.71
%75	MP	89.29	94.17	3.12	99.40
	$\alpha = 0.8$	76.83	95.50	6.56	99.45
	$\alpha = 0.5$	66.03	96.31	8.76	99.26
	$\alpha = 0.2$	54.70	97.77	10.90	99.14
%100	MP	90.03	100.00	2.98	100.00
	$\alpha = 0.8$	79.44	100.00	5.81	100.00
	$\alpha = 0.5$	72.15	100.00	7.26	100.00
	$\alpha = 0.2$	60.04	100.00	9.52	100.00

Table 1: Correctness percentages for various levels of homogeneity (HL) and s-norm decision (SD) types.

and a high value the social norms. Contrary to systems where the uploader controls the access of mutual content, PRINOR involves the co-owners in the final access control settings.

Considering the emergent *p-norms* according to Table 1, it is seen that *p-norm* assignments are almost always correct. They only depend on subsets of the decisions made by the same co-owners for the similar content types, and unless these agents change their behavior within their own co-owner groups, the *p-norm* based decisions would be the same as the decisions taken with the privacy decision mechanism. However, *p-norm* based decisions are usually a small part of all decisions, mostly due to sparsity of contents over different co-owner agent groups. *p-norms* require more mechanism based privacy decisions to emerge, and with emergent *s-norms*, *p-norms* build up slower than *s-norms*. Note that identifying *s-norms* requires a centralized location that holds the privacy decisions of the society. The OSN itself could provide this location and identification service. If there is no centralized location to enable identification of *s-norms*, PRINOR can still work with *p-norms*, as these are identified in a distributed manner and capture the norms of smaller groups of agents.

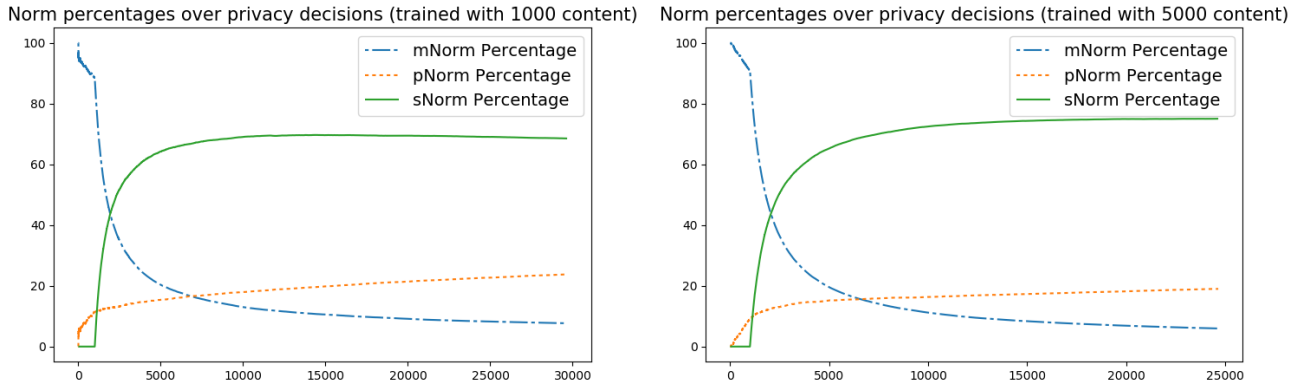


Figure 4: Percentage of norms over privacy decisions.

6 CASE STUDY FOR REAL-LIFE OSNS

Our simulation results show that with a given set of contextual features, privacy decisions for OSN content can be made successfully, both with *p-norms* and *s-norms*. However, real life OSNs pose further challenges. First, identifying contextual features of shared content is usually difficult, since users who share the content do not provide these properties. Thus, the contextual identification phase should rely on either the OSN provider or software agents that represent users. Second, some users would potentially have closer relationships and co-own more content than users with limited connections, necessitating the social network to reflect this. We tackle these with a case study by making use of two real-life data sets, namely *SNAP* [19] and *PicAlert* [39] and demonstrate the applicability of our approach.

PicAlert data set consists of images that have been annotated as private or public by study participants while *SNAP* data set contains friendship networks of Facebook users, including their bidirectional relationships, their circles, and anonymized personal features. Since our focus is on norm-based content privacy decisions by software agents on behalf of OSN users, our setup in this case study employs *PicAlert* for defining the contents and their privacy labels; and *SNAP* for defining agents and their network. We extract the contextual properties of *PicAlert* data set with an automated feature extraction tool, named *Clarifai* [38]. We assign four automatically generated tags to each *PicAlert* content as content descriptors and assign *SNAP* agents as co-owners.

Our setup in this case study is as follows: First, we generate all the possible *circles* between the agents in *SNAP* network. We define a circle as a relationship bond between multiple agents, where every two agents in the circle have an established relationship with each other. With this definition, each circle containing more than two agents would have subset circles, since all the subsets of a circle would still be a circle. Our second step is to pick content from *PicAlert* data set and to allocate the content to a circle. We run *Clarifai* API to get contextual tags for the image and assign the four most related tags of the image as content descriptors. Co-owners of a content are picked randomly from all possible circles. Second, we pick a number of content shared in the network and use them to generate m-norms for the co-owners of these shared content. We do this by considering the unique human decisions made for the

selected content (available through the data set) and matching the humans with the content co-owners. For example, if a human mostly has share decisions for a contextual tag, an m-norm with share action for the given tag is created for the matched co-owner. As a result of this, some agents acquire m-norms for possible future decisions, while the remaining agents do not establish a privacy understanding. This second category need to rely on emerging p-norms and s-norms to make decisions to protect their privacy.

The particular *SNAP* data set that we use has been extracted from Facebook and contains 347 users. These users have 975347 possible circles with two or more people. With the *PicAlert* data set, we generate content descriptors for 29864 content. We experiment with two setups of different initial content size. In the first setup, we extract 1000 content and generate m-norms to the co-owners of them. In the second setup, we increase the set to 5000 content. We repeat each setup five times, where different parts of the data set are picked as initial content. We display our results for the percentages of all norm types over incoming content in Figure 4 and correctness ratios after each 5000 privacy decisions in Figure 5.

Figure 4 shows the percentage of norms for access control decisions that has been taken in our system, from the first content considered to the last one. For both scenarios with different numbers of initial content, s-norms quickly start to emerge after the first 1000 decisions are made with either m-norms or p-norms. Around 5000 access decisions, s-norms converge to an amount of 70% with 1000 content used in training and 75% with 5000. Since these 5000 access decisions are for the content co-owned by one of all possible 975347 circles, the s-norms provide a majority of access control decisions without the need of another mechanism. According to Figure 4, we can conclude that even when most of the agents do not have any established m-norms, our approach is still able to form s-norms and p-norms to make privacy decisions. With a larger initial content set, s-norm ratio improves, because agents would usually access s-norms earlier, and would not have the need to make collaborative privacy decisions that eventually result in forming p-norms. Figure 5 plots the percentage of correct and incorrect p-norm and s-norm decisions after every 5000 access control decisions. When m-norms are formed from 5000 content, incorrect s-norms are only 1% of the decisions while p-norms are almost always correct. With 1000 content training scenario for m-norm formation, incorrect decision

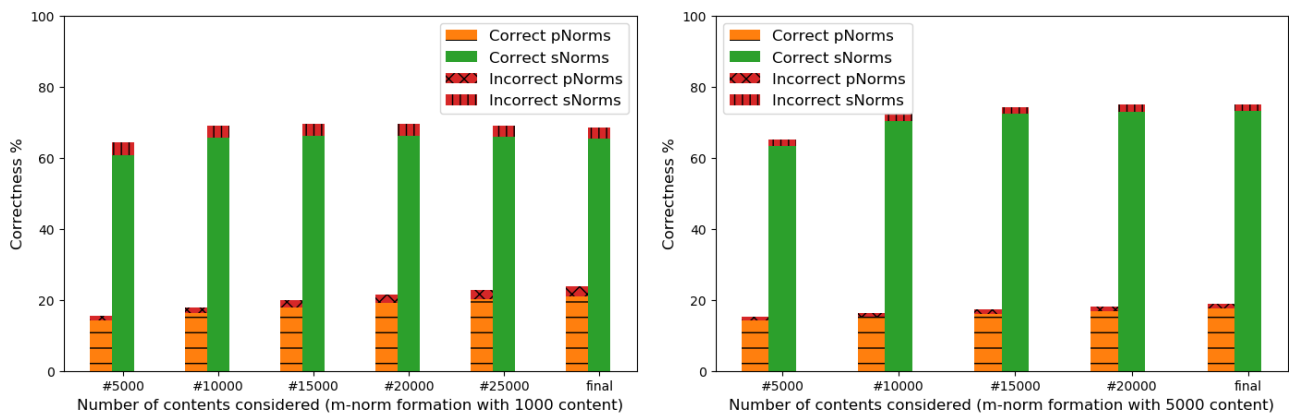


Figure 5: Correctness of p-norms and s-norms after each 5000 privacy decisions.

percentage gets slightly higher for both p-norms and s-norms, with s-norms reaching 3.5%, which can still be considered as a small part of the entire decisions. Thus, both setups achieve correct privacy decisions most of the time.

Performing a user study to verify whether the found social norms are in line with actual OSN users is beyond the scope of this paper. However, it is still useful to inspect a few of the norms to grasp the intuitions behind them. Combination of the following tags *baby*, *child*, *cute* and *little* generate social norms to deny access to users, as they are commonly seen as private. *Adult* tag, combined with tags such as *man*, *woman* or *person* again generate norms that deny access to other users. On the other hand, tags like *city*, *no person* and *nature* prescribe social norms that deem the content containing them as public, enabling access to other users. These examples indicate that the found social norms resonate with privacy understanding exhibited by many OSN users.

7 DISCUSSION AND CONCLUSION

Engineering privacy respecting methods for ubiquitous information systems has become crucial as the amount of online information is huge [13, 18, 28]. An important line of research focus on the specification and compliance of individual privacy preferences. Barth *et al.* [6] present a logic framework, formalizing aspects of contextual integrity and compliance with privacy norms. Barth *et al.* [7] study privacy for business processes, investigating if workflows would lead to data exposure or can verify that the privacy goals are achieved. Basin *et al.* [8] develop a monitoring tool to check policy compliance by employing first-order temporal logic for data relations.

There is a large body of work on access control in collaborative systems, especially online social networks. Hu *et al.* develop multiparty access control, where they develop a social network model, a multiparty policy specification scheme and a mechanism to enforce policies to resolve multiparty privacy conflicts [15]. Carminati *et al.* study a semantic web based framework to manage access control in OSNs by generating semantic policies [10]. The social network operates according to agreed system-level policies. Fong [12] pioneered the application of relationship-based access control mechanisms to collaborative systems, which initiated different lines of research. The

interoperability of relationship and role-based access control mechanisms is studied by Rizvi and Fong [24]. Mehregan and Fong [21] propose a policy negotiation mechanism for co-owned resources. These works provide feasible privacy resolution mechanisms for collaborative systems when policies are defined well. However, they require specification of policies for shared content offline, either by inference or with human expert involvement. Our work here, on the other hand, identifies the privacy norms that emerge in collaborative systems and makes them available to the users.

Norms have been studied in multiagent literature. Our previous work have investigated the idea of social norm emergence for OSNs [36]. However, it did not consider important aspects including prudential norms, aging of privacy decisions, or agent’s autonomy in choosing to follow norms. Calikli *et al.* [9] employ a social identity map for relationships of users and a set of social identity conflict rules to learn the privacy norms for social networks. Mashayekhi *et al.* [20] study norm emergence in traffic domain, where agents enter and leave and no known network structure among them exists. Ajmeri *et al.* [2] study norm emergence factoring in the context of the agents, taking in the sanctions into account. Our work is orthogonal to these work in the sense that we investigate norms in privacy, where agents cannot be required to follow them.

Such *et al.* perform an extensive, empirical evaluation to understand the dynamics of privacy with co-owned content [30]. They indicate that various co-owner type relations as well as different type of handling of violations might exist. Our work could serve as a solution to the problems identified there, as by identifying norms and applying them as they see fit, the agents can avoid privacy violations to take place. Thuraisingham *et al.* [33] tackle the privacy-awareness in handling data that is collected for business and marketing purposes and discusses design issues in achieving privacy-aware data management frameworks. Since only a small amount of privacy policies are known for that domain, our approach could help identify social norms.

8 LIMITATIONS AND FUTURE DIRECTIONS

PRINOR can generate societal norms both for the entire community and small sets of groups within; even after only a few privacy decisions have been taken. However, in its current state, it has some

limitations and room for improvements. First, current norm representations cannot express some of the interesting deontic concepts. Further, the interaction between norm types is limited. For example, agents do not update their moral norms after witnessing social norms. A more expressive representation of norms and their life cycles would enable the system to capture the established norms better, resulting in more successful collaborative access control decisions. Another limitation is that social norm emergence from the community behavior is reliant on OSN providers, while prudential norms for smaller groups can emerge with a distributed approach. Even though we prohibit OSN providers to obtain the entire privacy requirements of the users, a fully distributed approach can provide a more trustworthy mechanism to OSN users, free from possible tampering with the social norms by the providers. Our normative approach currently considers contextual properties of content according to automatically generated tags by tools, which may not include properties such as the location, time or the implicit information that can not be obtained without reasoning. This is another drawback we aim to tackle with our future work, where an ontology-based approach can enable an improved understanding of contexts in content descriptors so that social norms emerge accordingly with geographical, time-dependent, specific event related information and so on.

REFERENCES

- [1] 2018. General Data Protection Regulation. Available at: <https://ec.europa.eu/commission/priorities/justice-and-fundamental-rights/data-protection/2018-reform-eu-data-protection-rules>.
- [2] Nirav Ajmeri, Hui Guo, Pradeep K Murukannaiah, and Munindar P Singh. 2018. Robust Norm Emergence by Revealing and Reasoning about Context: Socially Intelligent Agents for Enhancing Privacy.. In *Proceedings of the International Joint Conference on AI (IJCAI)*, 22–34.
- [3] Natasha Alechina, Mehdi Dastani, and Brian Logan. 2012. Programming Norm-aware Agents. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2 (AAMAS '12)*, 1057–1064.
- [4] Md. Zulfikar Alom, Barbara Carminati, and Elena Ferrari. 2019. Helping Users Managing Context-Based Privacy Preferences, Elisa Bertino, Carl K. Chang, Peter Chen, Ernesto Damiani, Michael Goul, and Katsunori Oyama (Eds.). 100–107.
- [5] Leila Bahri, Barbara Carminati, and Elena Ferrari. 2018. Decentralized privacy preserving services for Online Social Networks. *Online Social Networks and Media* 6 (2018), 18–25.
- [6] A. Barth, A. Datta, J. C. Mitchell, and H. Nissenbaum. 2006. Privacy and contextual integrity: framework and applications. In *2006 IEEE Symposium on Security and Privacy (S P'06)*, 15 pp.–198.
- [7] A. Barth, J. Mitchell, A. Datta, and S. Sundaram. 2007. Privacy and Utility in Business Processes. In *20th IEEE Computer Security Foundations Symposium (CSF'07)*, 279–294.
- [8] David Basin, Matúš Harvan, Felix Klaedtke, and Eugen Zălinescu. 2012. MONPOLY: Monitoring Usage-Control Policies. In *Runtime Verification*. Springer Berlin Heidelberg, 360–364.
- [9] Gul Calikli, Mark Law, Arosha K. Bandara, Alessandra Russo, Luke Dickens, Blaine A. Price, Avelie Stuart, Mark Levine, and Bashar Nuseibeh. 2016. Privacy Dynamics: Learning Privacy Norms for Social Software. In *Proceedings of the 11th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS '16)*, ACM, 47–56.
- [10] Barbara Carminati, Elena Ferrari, Raymond Heatherly, Murat Kantarcioglu, and Bhavani Thuraisingham. 2009. A semantic web based framework for social network access control. In *Proceedings of the 14th ACM symposium on Access control models and technologies*, ACM, 177–186.
- [11] Francien Dechesne, Gennaro Di Tosto, Virginia Dignum, and Frank Dignum. 2013. No smoking here: values, norms and culture in multi-agent systems. *Artificial Intelligence and Law* 21, 1 (01 Mar 2013), 79–107.
- [12] Philip W.L. Fong. 2011. Relationship-based Access Control: Protection Model and Policy Language. In *Proceedings of the First ACM Conference on Data and Application Security and Privacy (CODASPY '11)*, ACM, 191–202.
- [13] Seda Gurses, Carmela Troncoso, and Claudia Diaz. 2011. Engineering Privacy by Design. *Computers, Privacy & Data Protection* (2011).
- [14] Chris Haynes, Michael Luck, Peter McBurney, Samhar Mahmoud, Tomas Vitek, and Simon Miles. 2017. Engineering the emergence of norms: a review. *The Knowledge Engineering Review* 32 (2017), e18.
- [15] Hongxin Hu, Gail-Joon Ahn, and Jan Jorgensen. 2013. Multiparty Access Control for Online Social Networks: Model and Mechanisms. *IEEE Transactions on Knowledge and Data Engineering* 25, 7 (July 2013), 1614–1627.
- [16] Dilara Kekulluoglu, Nadin Kokciyan, and Pinar Yolum. 2018. Preserving Privacy As Social Responsibility in Online Social Networks. *ACM Trans. Internet Technol.* 18, 4, Article 42 (April 2018), 22 pages.
- [17] Nadin Kökciyan, Nefise Yaglikci, and Pinar Yolum. 2017. An Argumentation Approach for Resolving Privacy Disputes in Online Social Networks. *ACM Trans. Internet Technol.* 17, 3, Article 27 (June 2017), 22 pages.
- [18] Marc Langheinrich. 2001. Privacy by Design — Principles of Privacy-Aware Ubiquitous Systems. In *UbiComp 2001: Ubiquitous Computing*, Gregory D. Abowd, Barry Brumitt, and Steven Shafer (Eds.), 273–291.
- [19] Jure Leskovec and Rok Sosič. 2016. SNAP: A General-Purpose Network Analysis and Graph-Mining Library. *ACM Trans. Intell. Syst. Technol.* 8, 1, Article Article 1 (July 2016), 20 pages.
- [20] Mehdi Mashayekhi, Hongying Du, George F List, and Munindar P Singh. 2016. Silk: A Simulation Study of Regulating Open Normative Multiagent Systems.. In *Proceedings of the International Joint Conference on AI (IJCAI)*, 373–379.
- [21] Pooya Mehregan and Philip W.L. Fong. 2016. Policy Negotiation for Co-Owned Resources in Relationship-Based Access Control. In *Proceedings of the 21st ACM on Symposium on Access Control Models and Technologies (SACMAT '16)*, Association for Computing Machinery, 125–136.
- [22] Javier Morales, Maite Lopez-Sanchez, Juan A. Rodriguez-Aguilar, Michael Wooldridge, and Wamberto Vasconcelos. 2013. Automated Synthesis of Normative Systems. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems (AAMAS '13)*, 483–490.
- [23] Federica Paci, Anna Squicciarini, and Nicola Zannone. 2018. Survey on Access Control for Community-Centered Collaborative Systems. *Comput. Surveys* 51, 1, Article 6 (Jan. 2018), 38 pages.
- [24] Syed Zain R. Rizvi and Philip W.L. Fong. 2016. Interoperability of Relationship- and Role-Based Access Control. In *Proceedings of the Sixth ACM Conference on Data and Application Security and Privacy (CODASPY '16)*, 231–242.
- [25] Ravi S. Sandhu, Edward J. Coyne, Hal L. Feinstein, and Charles E. Youman. 1996. Role-Based Access Control Models. *IEEE Computer* 29, 2 (1996), 38–47.
- [26] Bastin Tony Roy Savarimuthu and Stephen Crane. 2011. Norm Creation, Spreading and Emergence: A Survey of Simulation Models of Norms in Multi-agent Systems. *Multiagent Grid Syst.* 7, 1 (Jan. 2011), 21–54.
- [27] Sandip Sen and Stéphane Airiau. 2007. Emergence of norms through social learning.. In *Proceedings of the International Joint Conference on AI (IJCAI)*, Vol. 1507, 1512.
- [28] Sarah Spiekermann and Lorrie Faith Cranor. 2009. Engineering Privacy. *IEEE Trans. Softw. Eng.* 35, 1 (Jan. 2009), 67–82.
- [29] Anna C. Squicciarini, Mohamed Shehab, and Federica Paci. 2009. Collective Privacy Management in Social Networks. In *Proceedings of the 18th International Conference on World Wide Web (WWW '09)*, ACM, 521–530.
- [30] Jose M Such, Joel Porter, Sören Preibusch, and Adam Joinson. 2017. Photo privacy conflicts in social media: A large-scale empirical study. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, ACM, 3821–3832.
- [31] Jose M. Such and Michael Rovatsos. 2016. Privacy Policy Negotiation in Social Media. *ACM Transactions on Autonomous and Adaptive Systems* 11, 1, Article 4 (Feb. 2016), 29 pages.
- [32] Pang-Ning Tan, Michael Steinbach, and Vipin Kumar. 2005. *Introduction to Data Mining, (First Edition)*. Addison-Wesley Longman Publishing Co., Inc.
- [33] Bhavani Thuraisingham, Murat Kantarcioglu, Elisa Bertino, Jonathan Z. Bakhdash, and Maribel Fernandez. 2018. Towards a Privacy-Aware Quantified Self Data Management Framework. In *Proceedings of the 23rd ACM on Symposium on Access Control Models and Technologies (SACMAT '18)*, 173–184.
- [34] Raimo Tuomela. 1995. *The Importance of Us: A Philosophical Study of Basic Social Norms*. Stanford University Press.
- [35] Onuralp Ulusoy and Pinar Yolum. 2018. PANO: Privacy Auctioning for Online Social Networks. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '18)*, 2103–2105.
- [36] Onuralp Ulusoy and Pinar Yolum. 2019. Emergent Privacy Norms for Collaborative Systems. In *PRIMA 2019: Principles and Practice of Multi-Agent Systems*. Springer International Publishing, Cham, 514–522.
- [37] Piotr A. Woźniak, Edward J. Górzelańczyk, and Janusz A. Murakowski. 1995. Two components of long-term memory. *Acta neurobiologiae experimentalis* 55, 4 (1995), 301–305.
- [38] Matthew D. Zeiler and Rob Fergus. 2014. Visualizing and Understanding Convolutional Networks. In *Computer Vision – ECCV 2014*, David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (Eds.), Cham, 818–833.
- [39] Sergej Zerr, Stefan Siersdorfer, and Jonathon Hare. 2012. PicAlert! A System for Privacy-Aware Image Classification and Retrieval. In *Proceedings of the 21st ACM International Conference on Information and Knowledge Management (CIKM '12)*, Association for Computing Machinery, 2710–2712.