Emergent Privacy Norms for Collaborative Systems

Onuralp Ulusoy(✉) and Pınar Yolum

Utrecht University, Utrecht, The Netherlands
{o.ulusoy,p.yolum}@uu.nl

Abstract. Managing privacy of users in online systems is a major aspect of cyber-security. Typical approaches to privacy are concerned with giving users options of informed consent, wherein users define their private data, how they want them to be used, and so on. However, in collaborative systems, such as online social networks, managing privacy exhibits problems beyond traditional consent, since a content being shared (such as a group picture or a multi-party business contract) might belong to more than a single entity, with different privacy policies. Recent approaches to preserve privacy in such settings rely on multiagent agreement technologies, which require a new decision to be formed for every content that will be shared, making them difficult to scale for real life applications. Accordingly, this paper proposes a normative approach for maintaining privacy in collaborative systems that do not require a decision to be formulated from scratch for each content. Instead, the system generates social norms based on previous decisions. The agents are free to follow the social norms as well as their own privacy policies. We show over multiagent simulations that our approach extracts social norms successfully and enables successful privacy decisions to be taken.

Keywords: Privacy · Multiagent systems · Norm emergence

1 Introduction

Collaborative systems, such as online social networks (OSNs), contain tremendous amount of content. These content, being shared by the OSN users, can be related to multiple people, as in the example of a group picture. However, these kind of content also might contain private information of people, either explicitly or implicitly. Hence, the decision of sharing or not sharing a content should be decided collaboratively by the users who are affected by it.

Collaborative privacy management mechanisms aim to resolve the conflicts in such cases. Finding a suitable resolution is usually a challenging task, since satisfying privacy protection constraints of some users might result in not sharing content that other users wanted to share, which is also undesirable by OSN providers since it would cause fewer content to be shared in the network. Multiagent agreement technologies, such as argumentation [6], negotiation [5,14] or

M. Baldoni et al. (Eds.): PRIMA 2019, LNAI 11873, pp. 514–522, 2019.
https://doi.org/10.1007/978-3-030-33792-6_36

auctions [13,17] have been successfully used for resolving privacy disputes. But, these approaches have two major drawbacks. First, they require each agent to actively participate in the decision making mechanism whenever a content is relevant to them. This is a large overhead for systems where a large set of content is shared regularly. For example, if nude pictures are never shared, it is redundant to deliberate on an incoming nude picture. Second, they ignore the relations and background knowledge available to the agents that are involved in the system as well as the values of the society that cannot be reflected with individual decisions of the agents. For example, a group of friends might share their Friday outing pictures regularly without a need to come to a mutual agreement on whether this is private or not every single time. To overcome the listed deficiencies, we propose to use normative multiagent systems where privacy decisions are taken based on the norms that are generated from the privacy decisions in the system. It is well known that human societies are guided and controlled by the norms [11]. Since privacy decisions over OSNs are correlated with the society behavior; extracting the norms from previous privacy actions could make them useful for future privacy decisions. If the norms are not applicable in a given situation (based on the content type or the individuals involved), the system still employs an existing collaborative privacy management mechanism to make a decision.

2 Privacy Norms

A typical collaborative system is online social networks, where users share content about themselves as well as others. A generic OSN consists of three main elements: users, a set of relations between the OSN users and content that is shared within the social networks of the users. Users can have privacy expectations that can vary based on the type of content as well as the other users to whom the content is exposed. 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. In OSNs, it is common for a content to contain private information of users, either explicitly (e.g., geotagging or name tagging) or implicitly (e.g., finding about the location from the visible objects in the background, content's sharing time implying extra information about the co-owners and so on). When a single content contains private information of more than one user, conflicts might occur; some users wanting to share a content in the OSN while others want to share it with only a limited number of users, or even not share it at all. This requires a decision mechanism to be in place, so that for a new content the system can reach a *privacy decision* as to how it will be shared and with whom.

To reach privacy decision effectively and efficiently, we design a normative multiagent system, where privacy expectation of users for sharing content are being managed by software agents [6]. We represent a content with (i) a content type matrix which stores the contextual properties of the content, such as holiday, work, and so on and (ii) a set of co-owner agents whose privacy is possibly being affected by the content and thus should have a say about content's privacy decision. The system contains norms to capture the privacy preferences.

Informally, privacy norms capture the acceptable behaviour for sharing a particular type of content with a particular set of users. The acceptable behavior need not always be understood as the expected behavior of the majority. If a community is formed by privacy aware agents, an agent's request to not share a content might be complied by the others who initially wanted the content to be shared. If such an example occurs frequently, it can emerge as a norm and can be enforced to future decisions with the same context, where agents are not that privacy aware. In literature, privacy related access control mechanisms either enforces the strongest action (e.g., majority action) or rule based privacy decisions (e.g., deny overriding other actions). Minority protecting norms can aid a mechanism to differ from previous works in this sense, which could be beneficial for the goal of more satisfactory privacy decisions.

We adopt Tuomela's categorization of norms; where personal norms contain *m-norms* (i.e. moral norms) and *p-norms* (i.e. prudential norms), while social norms contain *r-norms* (i.e. rule norms) and *s-norms* (i.e. social norms) [16]. We formally represent them similar to existing formalisms [2,9], such that a set of preconditions determine the activation of a sharing action to be taken. Since our focus is more on the emergence of norms rather than their violation, we do not include norm sanctions explicitly. Thus, we employ *s-norms* as social norms, while *m-norms* contain all privacy requirements of individual agents.

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<rType, cClass, act{share,noShare}>*, where *rType* is the main relationship context between the co-owners for a content, *cClass* is the specific class of the similar contents and *act* is the assigned action of the norm, which could be either sharing or not sharing the content. *s-norms* emerge depending on the previous collaborative decisions within the OSN. We employ *rType* since *s-norms* are generated according to an overview of the societal decisions and *cClass* because the norms pertain to the generic behavior of the society.

m-norms are based on individuals' own preferences (i.e., understanding of what is right to do). An agent might prefer not to share a content that it thinks is offensive to others. We represent *m-norms* as 3-tuple as well: *m<rType, cType, act{share,noShare}>*, where *rType* is the main relationship type, *cType* is the major content type, and *act* is the action to take when these conditions are satisfied.

3 Normative Privacy Decision

The agents' personal privacy expectations are represented as *m-norms* and stored in a personal *m-norm* base, which can only be changed or updated by the agent itself. *s-norm* base contains the social norms, which emerge based on the privacy decisions of the individual agents. There is a single *s-norm* base in the system. The normative decision mechanism process progresses with every incoming content. Initially, agents only have *m-norms*. *s-norms* emerge over time based on the actions of the agents. All types of norms have a lifecycle, where they are created, updated, or removed from the respective norm base. Our approach enables agents to make a privacy decision based on the norms in the system first and if that is not possible reverts to a collaborative decision mechanism.

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 two types of norms are in effect, there can easily be conflicts among these 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.

In this work, we assume *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 *s-norm* base to see if a norm matching with the content type matrix exists. If so, it is applied without triggering the collaborative privacy decision mechanism within the system. It might be the case that none of the norms in the social norm base are applicable to make a decision. In that case, the decision mechanism is triggered and the final decision is made according to the collaborative privacy mechanism, which makes use of *m-norm* bases of all co-owner agents. In the latter case, the outcome of the mechanism also updates the *s-norm* base of the OSN, where new possible norms can be formed for future co-owned content.

It is crucial for norms to be identified and managed accurately. *m-norms* are private to each user and thus managed individually. The management of *s-norms* are more challenging because they emerge and die based on the users' interactions. *s-norms* reflect the privacy choices of the society as a whole and emerge based on the previous privacy decisions that are taken by the users on a given content type. For example, if in many occasions, the users that are colleagues do not share content about their holiday, this can emerge as an *s-norm* in the society. Since the OSN provider has access to all the privacy decisions in the system, the lifecycle of an *s-norm* can be managed by the OSN provider. Given the previous privacy decision, how can an OSN generate *s-norms*? To achieve this, we develop Algorithm 1. Algorithm 1 generates the norms from decisions using the intuition that we place all content over a multidimensional space according to their content type matrix dimensions and the relationship type of the co-owner agents. This space contains all the decisions considering its various properties as dimensions. Next, we cluster this space such that each cluster contains content that have similar attributes. Then, the clusters can be assigned as *s-norm* classes, and can be checked for normative behavior; i.e., qualified majority of the content in the same cluster result in the same type of sharing action. We call this type of clusters as *normative clusters*. To ensure that generated norms are still in effect, the algorithm is run periodically and the *s-norm* base is updated accordingly.

OSNs contain a tremendous amount of content; thus, continuously clustering the content space would require massive computing power. Since dimensions of *s-norms* are interrelated (e.g., content type matrix dimensions), neighbor clusters could have similar normative behavior and contextual properties. This enables us to relax the problem precision requirements, allowing us to not necessarily find the optimum solution for a few borderline decisions, but place them in one of the closest cluster. This approach also turns up beneficiary for our goal, since a borderline privacy decision between two clusters would mean that the decision is in a similar distance from both clusters and not strongly related with a single one. On the contrary with the privacy decisions in the center of the clusters, these kind of privacy decisions could belong to both of the clusters with weak ties. To achieve this, we employ *k-means* algorithm to cluster content and then check 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 [15]. Note that the size of the clusters is important as they affect the number of clusters. Having few clusters with large amounts of content would result in not discovering normative behavior, while clustering with fewer number of content in each cluster would result in increased complexity. We address this by starting with a small number of clusters and increasing the number of clusters iteratively. The iteration for elements in a cluster ends when the threshold for minimum number of agents that a cluster can contain is reached, or a normative behavior is already found.

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
Output: cList, a set of clusters generated from pDec
while pDec not empty do
    tempcList = k-Means(mk,pDec)
    foreach cluster in tempcList do
        isNormative = checkNormative(cluster)
        if (isNormative = true or size(cluster) < t) then
            add(cluster,cList)
            foreach item in cluster do
                remove(item,pDec)
    mk += 1
return cList
```

Algorithm 1 takes the minimum cluster count parameter (mk), the minimum size threshold parameter for a single cluster (t) and all the previous privacy mechanism based decisions ($pDec$) as input. It then assigns all the items in $pDec$ to a cluster in the output cluster list ($cList$). In each iteration, a temporary list of clusters are assigned with k-means algorithm, where all items in $pDec$ are clustered and the number of clusters are given as *mk*. In line 3, a for loop begins, which checks the temporary cluster assignments, and determines if the cluster shows a normative behavior, or the size of the cluster is below t value. If one of these conditions are satisfied for a temporary cluster, the cluster is added to *cList* in line 6 and all the items of the cluster are removed from *pDec*, ending an 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. When a new content comes, agents find the most similar cluster, according to its content type matrix and the relation between the co-owners. If this is a normative cluster, then agents can decide according to the related normative action.

4 Evaluation

Our main goal is to reduce the necessity of applying collaborative privacy management algorithms by identifying the emergent norms within OSNs. We study the emergence of norms through multiagent simulations in an environment we developed in Java. Each agent in the simulation represents a user. The users, and thus the agents are related to each other through one relationship. Each agent has a set of *m-norms* that are generated automatically. Each content in the OSN is related to a set of contexts with varying levels and is thus represented with a content type matrix. In real life, this information would come from the features or tags of the content. Here, we assume that the matrix, where major content type categories are predefined is given. 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. Each content has a set of co-owners, which are the agents with private information represented in the content.

We include 100 agents and 10000 contents for each of our simulations, where each content is randomly assigned to 2 to 5 co-owners, and a 4 dimensional content type matrix, while each dimension has 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 *s-norms*. Each simulation follows the flow in Sect. 2. For simplicity, we employ majority voting as the collaborative privacy mechanism in our evaluations.

Societal normative behavior for privacy emerges when a set of agents have a similar idea of privacy; e.g., prefer to share similar content. If agents have totally different views about privacy, we do not expect norms to emerge. On the other hand, if all agents share the same idea of privacy, then there would be a few norms that could govern the entire population. We expect many populations to stand between these two cases. To account for this, we introduce a variable to capture the homogeneity of a given society. If the homogeneity of the society is 0%, then all the agents in the population can have different privacy choices. We ensure this by allowing them to randomly make a choice about sharing or not sharing

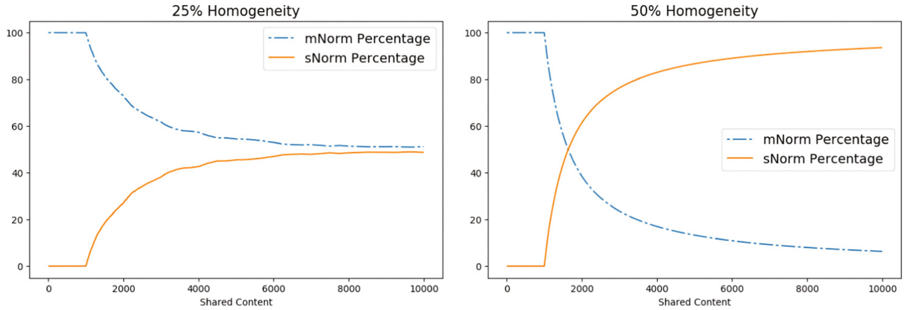

Fig. 1. Percentage of norm types over different levels of homogeneity.

a given content. We run the simulations with several levels of homogeneity. The simulation starts forming *s-norms* using Algorithm 1 after 1000th content shared in the OSN and reruns it after every 250 content for updating the *s-norm* base of the OSN. For each homogeneity level, we run 5 simulations and measure the percentage of our norm types over the number of content shared within the environment. This shows the necessary level of homogeneity for social norms to emerge.

Figure 1 plots the percentage of decisions that are taken by *m-norms* and *s-norms* as new content is introduced to the system for populations with two homogeneity levels. We omit other homogeneity levels for brevity as these two levels are sufficient enough to capture the trend. The plots show that if a quarter of the community shows homogeneous behavior, almost half of the decisions can be made according to social norms, reducing the need to use a decision mechanism to half. When half of the community behaves the same, only less than 10% of the decisions require a collaborative decision mechanism, and social norms can be decisive for more than 90% of privacy decisions.

5 Discussion

We have investigated how a normative approach can cope with privacy protection in a multiagent system that contains collaboration, cooperation and competition aspects for the agents at the same time. We apply our method thoroughly for the OSN domain, and evaluate it over multiagent simulations in terms of to what extent the privacy issues can be resolved with norms and their correctness in their resolutions.

Engineering privacy in ubiquitous information systems has become a research interest after millennium, mostly because internet becoming a part of a daily life with OSNs, smart devices etc. and causing massive amount of private information to be accessible by the others. Langheinrich [7] investigates the open issues for privacy-respecting approaches for ubiquitous computing. Spiekermann and Cranor [12] and Gurses [4] study the grounds of engineering privacy, explaining how information related domains can be designed and employ privacy-preserving methods. Paci *et al.* [10] provide an extensive survey for literature about access control over community centric collaborative systems; laying down the key issues and giving a roadmap for future challenges. Bahri *et al.* [3] show the challenges of preserving privacy over decentralized OSNs, and provides a review of previous work done for overcoming these challenges. These studies all show that privacy is an important aspect of information systems and should be tackled to prevent violations.

Sen and Airiau [11] pioneered the work of norm emergence, where they show that even when the population size and heterogeneity vary, social norms can emerge. Mashayekhi *et al.* [8] investigate norm emergence in traffic domain, where agents enter and leave and no known network structure among them exists. Ajmeri *et al.* [1] study norm emergence factoring in the context of the agents, taking in the sanctions into account. Our findings here show that for privacy norms to emerge, it is enough for the population to have a low level of homogeneity in how they perceive privacy.

An interesting direction is to study norm emergence when agents' privacy expectations can change over time as they learn new facts, face new norms or as their relationships change. Another direction is to enable agents to judge the social norms based on their own privacy values. We also aim to implement a mechanism of forgetting for norms, where social norms can fade over time. These would bring us closer to accommodating groups with different privacy norms to coexist in a society.

References

1. Ajmeri, N., Guo, H., Murukannaiah, P.K., Singh, M.P.: 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) (2018)
2. Alechina, N., Dastani, M., Logan, B.: Programming norm-aware agents. In: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems. AAMAS 2012, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, vol. 2, pp. 1057–1064 (2012)
3. Bahri, L., Carminati, B., Ferrari, E.: Decentralized privacy preserving services for online social networks. Online Soc. Netw. Media **6**, 18–25 (2018)
4. Gurses, S., Troncoso, C., Diaz, C.: Engineering privacy by design. Comput. Priv. Data Prot. **14**(3), 25 (2011)
5. Kekulluoglu, D., Kokciyan, N., Yolum, P.: Preserving privacy as social responsibility in online social networks. ACM Trans. Internet Technol. **18**(4), 42:1–42:22 (2018)
6. Kökciyan, N., Yaglikci, N., Yolum, P.: An argumentation approach for resolving privacy disputes in online social networks. ACM Trans. Internet Technol. **17**(3), 27:1–27:22 (2017)
7. Langheinrich, M.: Privacy by design — principles of privacy-aware ubiquitous systems. In: Abowd, G.D., Brumitt, B., Shafer, S. (eds.) UbiComp 2001. LNCS, vol. 2201, pp. 273–291. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-45427-6_23

8. Mashayekhi, M., Du, H., List, G.F., Singh, M.P.: Silk: a simulation study of regulating open normative multiagent systems. In: Proceedings of the International Joint Conference on AI (IJCAI), pp. 373–379 (2016)
9. Morales, J., Lopez-Sanchez, M., Rodriguez-Aguilar, J.A., Wooldridge, M., Vasconcelos, W.: Automated synthesis of normative systems. In: Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems. AAMAS 2013, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, pp. 483–490 (2013)
10. Paci, F., Squicciarini, A., Zannone, N.: Survey on access control for community-centered collaborative systems. ACM Comput. Surv. **51**(1), 6:1–6:38 (2018)
11. Sen, S., Airiau, S.: Emergence of norms through social learning. In: Proceedings of the International Joint Conference on AI (IJCAI), vol. 1507, p. 1512 (2007)
12. Spiekermann, S., Cranor, L.F.: Engineering privacy. IEEE Trans. Softw. Eng. **35**(1), 67–82 (2009)
13. Squicciarini, A.C., Shehab, M., Paci, F.: Collective privacy management in social networks. In: Proceedings of the 18th International Conference on World Wide Web WWW 2009, pp. 521–530. ACM, New York (2009)
14. Such, J.M., Rovatsos, M.: Privacy policy negotiation in social media. ACM Trans. Auton. Adapt. Syst. **11**(1), 4:1–4:29 (2016)
15. Tan, P.N., Steinbach, M., Kumar, V.: Introduction to Data Mining, 1st edn. Addison-Wesley Longman Publishing Co., Inc., Boston (2005)
16. Tuomela, R.: The Importance of Us: A Philosophical Study of Basic Social Norms. Stanford University Press, Stanford (1995)
17. Ulusoy, O., Yolum, P.: Pano: privacy auctioning for online social networks. In: Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems. AAMAS 2018, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, pp. 2103–2105 (2018)