PANOLA: A Personal Assistant for Supporting Users in Preserving Privacy

ONURALP ULUSOY and PINAR YOLUM, Utrecht University, The Netherlands

Privacy is the right of individuals to keep personal information to themselves. When individuals use online systems, they should be given the right to decide what information they would like to share and what to keep private. When a piece of information pertains only to a single individual, preserving privacy is possible by providing the right access options to the user. However, when a piece of information pertains to multiple individuals, such as a picture of a group of friends or a collaboratively edited document, deciding how to share this information and with whom is challenging. The problem becomes more difficult when the individuals who are affected by the information have different, possibly conflicting privacy constraints. Resolving this problem requires a mechanism that takes into account the relevant individuals' concerns to decide on the privacy configuration of information. Because these decisions need to be made frequently (i.e., per each piece of shared content), the mechanism should be automated. This article presents a personal assistant to help end-users with managing the privacy of their content. When some content that belongs to multiple users is about to be shared, the personal assistants of the users employ an auction-based privacy mechanism to regulate the privacy of the content. To do so, each personal assistant learns the preferences of its user over time and produces bids accordingly. Our proposed personal assistant is capable of assisting users with different personas and thus ensures that people benefit from it as they need it. Our evaluations over multiagent simulations with online social network content show that our proposed personal assistant enables privacy-respecting content sharing.

CCS Concepts: • **Security and privacy** → **Privacy protections;**

Additional Key Words and Phrases: Autonomous agents, online social networks, privacy, reinforcement learning

ACM Reference format:
Onuralp Ulusoy and Pınar Yolum. 2021. PANOLA: A Personal Assistant for Supporting Users in Preserving Privacy. *ACM Trans. Internet Technol.* 22, 1, Article 27 (September 2021), 32 pages.
https://doi.org/10.1145/3471187

1 INTRODUCTION

Many of the recent software systems are built on the idea of collaborative computing, where multiple users present, manipulate, and, as a result, manage shared content. While previously multiple

This research was (partially) funded by the Hybrid Intelligence Center, a 10-year programme funded by the Dutch Ministry of Education, Culture and Science through the Netherlands Organisation for Scientific Research, https://hybrid-intelligence-centre.nl, grant number 024.004.022.

Authors' address: O. Ulusoy and P. Yolum, Department of Information and Computing Sciences, Utrecht University, Princetonplein 5, 3584 CC, Utrecht, the Netherlands; emails: {o.ulusoy, p.yolum}@uu.nl.

1533-5399/2021/09-ART27 $15.00
https://doi.org/10.1145/3471187

users would only access their own data, such as e-commerce systems or banking systems, now the information is being accessed, edited, and served to others by many. Consider an **online social network (OSN)**, where a user can share pictures that include other users, who are many times able to tag themselves or others, comment on it, and even reshare it with others. Or, consider an IoT system, in which one security camera would like to share footage of a setting to guarantee security for all people, while one individual would prefer to keep the location of him- or herself secret. In both cases, the content in question relates to multiple entities who have different privacy concerns or expectations from each other. Even though the content is meant to be shared by a single entity, the content is related to more than just the uploader and hence is actually *co-owned* by others [19, 36].

Many times, the co-owned content is shared without an explicit consent from all of the co-owners. Even a seemingly unimportant piece of content might be considered private by one of the co-owners, and its sharing might have drastic effects. Ideally, when co-owners have different privacy constraints, the entities should be given the means to make a decision as to whether they wish to share the content. Reaching such a decision is usually not an easy task, since people's privacy requirements can easily be in conflict [34]. In real life, such decisions require time and effort as individuals interact to reach decisions. However, current systems enable only the uploader to set privacy settings while publishing content but do not allow co-owners to state their constraints. As a result, existing systems cannot provide collaborative solutions, and individuals are left to resolve conflicts via offline methods (e.g., personal communication) [25] or, in most cases, ignore others' privacy requirements and cause voluntary or involuntary privacy violations.

Ideally, systems should provide privacy management mechanisms to regulate how content will be shared. Recently, multiagent agreement techniques were used to address collaborative privacy management. Kekulluoglu et al. [20] and Such and Rovatsos [36] propose negotiation-based approaches that enable users to reach a consensus on how to share content. Kökciyan et al. [23] use argumentation to enable one user to persuade the other into sharing with his or her own privacy constraints. These approaches have been successful but can only be used when the entities can reason on the users' privacy policies and communicate with others intensively. Moreover, the entities in these systems follow predefined rules but do not learn to preserve their users' privacy more over time. Alternatively, Squicciarini et al. [32] propose a model where users enter auctions for deciding on a policy that requires collaborative management over content. Each user creates bids based on how much he or she wants to see some content public or private. In that approach, users earn points by publishing content and tagging people that are related to content. These points are used in an auction, where users spend their currencies to convince other users to accept their policy, based on the Clarke Tax mechanism [9, 14]. Ulusoy and Yolum [39, 41] have extended this mechanism into a system called PANO to ensure that the users cannot abuse the system. This is done by enforcing that the obtained points can be used only in the same groups and that a given bid can never go beyond a predefined maximum.

In addition to having a useful mechanism, it is important that users participate in the mechanisms to act based on their preferences. There are two difficulties with user participation in these mechanisms. First, many existing works show that users themselves do not usually know their privacy constraints, let alone evaluate the importance of contextual properties for privacy [1, 13]. Thus, when users take part in the mechanism, they might not participate in the way that would benefit them the most. Second, since the amount of content in OSNs is large, participating in such mechanisms for each type of content is time-consuming for many users. Thus, it is not realistic to assume that the users will take part in these mechanisms for all content shared.

To address these issues, we advocate a distributed approach, where each user in the system is represented by a personal assistant, which is a software agent that can perceive, reason, and act on behalf of its user. These personal assistants need to understand how they can help their users, learn their preferences over time, and perform the users' tasks in the mechanism on their behalf. First, the design of such a personal assistant needs to take into account two properties of the users: the privacy valuations of users for different types of content and the valuations of users for conforming with decisions of the groups of which they are part. These are important because both of these influence how a user would participate in a mechanism. For example, for a given piece of content, if the value of the content is high, the user might prefer to do whatever it takes to preserve its privacy. For a different piece of content, the user might prefer to cooperate with the rest of the group. Second, the design of a personal assistant needs to take into account the details of the mechanisms in place. The personal assistant participating in a negotiation would conduct reasoning different from one participating in an auction. In a similar vein, the personal assistant would need to learn different aspects of the mechanism to fulfill its task. For example, for a negotiation, the personal assistant might learn to formulate better counter offers, while for an auction, it would learn to generate correct bids.

Learning has been used in the context of privacy before, mostly to enable agents to classify whether a user would consider the content in question private or not [15, 30]. These approaches make use of the previous interactions of the user with the system to employ various supervised learning algorithms as well as information retrieval techniques to infer the privacy of content. However, the learning problem posed here has characteristics different from the problem that has been studied in the literature. First, what needs to be learned is not whether some content is private or not, but what the agent would bid to share or not to share the content. The bid would be affected by what the agent has shared before, whether that led to a beneficial outcome, what the user's valuation of the content initially was, and whether the user conforms to the group he or she is in. Second, existing learning algorithms for privacy consider a single user's point of view, but here the privacy has to be considered in a group, since the content to be shared is co-owned. Hence, other users' actions influence the outcome of a privacy decision. This creates the need to learn in the context of a given group of individuals. We tackle this learning problem with the use of *reinforcement learning* so that the agents can interpret the overall privacy decisions to adjust how they formulate their bids.

This article describes **Privacy Auctioning Learning Agent (PANOLA)**, which acts as a personal assistant to users in situations where a piece of co-owned content is being shared. For decision making, PANOLA uses PANO (see Section 2.1), which is robust and can thus accommodate a large number of decisions to be taken. PANOLA can make use of user input as an initial point to bid but then learns to adjust its bidding strategy over time. Our previous work identified possible factors that might be important for realizing reinforcement learning for generating bids (e.g., capturing bid ranges) and showed the effects of some of these factors in realistic settings (e.g., having unlimited budget for auctions [40]). Using these results, here, we develop a refined model to realize reinforcement learning, show how it can be used for decision making, and study in detail how PANOLA can be helpful for users in preserving privacy.

While helping users, there are two important criteria that need to be respected. The first is the extent to which PANOLA can help different types of users. It is well known that users can vary in their expertise in handling privacy. The personal assistant that we develop should be able to help users with different levels of knowledge and motivation in thinking about privacy. The second is that through the personal assistant, we would like to enable all users to have fair use of the system. That is, because the content in question is owned by many individuals, possibly with conflicting privacy preferences, it might not be possible to preserve every user's privacy for all content. When this is the case, it should be ensured that no user is left at a disadvantage, such that always the same individuals' privacy is being preserved while others' privacy is being violated.

Table 1. Four User Bids for Sharing an Image

Users	Not Share	Limited Share	Share
Alice	3	5	0
Bob	15	2	0
Carol	5	8	5
Dave	2	6	18

The rest of this article is organized as follows: Section 2 explains the necessary background on the PANO auctions and common personas in OSNs. It also introduces our running example. Section 3 describes how PANOLA works, with a focus on its learning. Section 4 explains our experimental setup and answers our research questions through multiagent simulations. Finally, Section 5 discusses our work in relation to existing methods in the literature and gives pointers for future work.

2 TECHNICAL BACKGROUND

We design PANOLA in the context of an auction decision-making mechanism, namely PANO, and serve different types of users that can be defined using *privacy personas*. We provide the technical background on these here.

2.1 PANO: Agent-Based Privacy Auctioning

PANO is an agent-based privacy decision system, where agents employ auctioning mechanisms to reach decisions on privacy conflicts [41]. PANO employs an extended version of the Clarke Tax mechanism as an underlying mechanism. The Clarke Tax mechanism [9] provides an auction mechanism, where participants bid for different possible outcomes in the environment. The outcome that receives the highest total bids from the participants wins and is carried out. Different from an English auction, participants who aid in the winning outcome to be chosen, i.e., who bid toward it, are taxed according to the value they put on it. This is calculated by subtracting the bid values of every single user from the overall values of the outcomes. If the subtraction of a single user's bid changes the overall decision, it shows that the user's bid on this outcome had a *decisive* value. Thus, the user is taxed with the difference of the actual outcome's score and the score of the outcome to be taken if that user were not present in the auction.

In the context of collaborative privacy, the Clarke Tax mechanism is used to decide on how content is going to be shared. PANO can work with as many outcomes as needed in a given context. Consider the example below with three possible outcomes: *not share*, *limited share*, and *share*. *Not share* depicts an outcome where the content is kept private and cannot be accessed by anyone; *limited share* is an outcome where the content is shared with a specific group of people, such as *colleagues*; and *share* means it can be accessed by everyone.

Table 1 shows an example of bidding for four users deciding how to share some content. Users decide based on their own value of the three outcomes. According to Table 1, it can be seen that Bob values the *not share* outcome more than the others, since its bid value for *not share* is 15, whereas for *limited share* it is 2 and for *share* 0. On the other hand, Dave values the *share* outcome the most as his bid value is 18. According to the bidding of all users, the Clarke Tax auction mechanism decides on which outcome will be in effect. Based on the bids from Table 1, the *not share* outcome receives a total of 25, whereas *limited share* receives a total of 21 and *share* a total of 23 points. Therefore, the *not share* outcome is chosen.

Table 2 shows the resulting decision and applies the taxes according to the bidding in Table 1. Each user who bids for the decisive outcome needs to be taxed. If Alice or Bob had not voted for *not*

Table 2. Clarke Tax Mechanism Example—Decision and Taxes

Values	Not Share	Limited Share	Share	Taxes
Overall	**25**	21	23	
Without Alice	22	16	**23**	1
Without Bob	10	19	**23**	13
Without Carol	**20**	13	18	0
Without Dave	**23**	15	5	0

share, the winning outcome would have been *share* instead. Hence, both Alice and Bob are taxed. When the score of Alice is subtracted from the overall score, the decision of sharing the content receives the maximum score by 23, while not sharing gets a score of 22. This causes Alice to be taxed with a score of 1. As mentioned above, Bob bid a greater value for the *not share* outcome, and its hypothetical absence from the auction also causes the resulting decision to change. Since the differences of the outcomes are much higher in Bob's case (i.e., 13, obtained from the subtraction of *share* and *not share* scores), Bob is taxed with a greater value. It is important to note that, although the user is taxed, the outcome he values the most is decided on and the content is not shared. Good evaluation and truthfulness for bidding are crucial. For example, if Bob bid for not sharing with a rather big value, even though the decision was not that important to him, he would have paid the bid amount plus a great amount of tax. On the contrary, if he had bid for much less, then his privacy might have been violated. Hence, it is important to be able to create bids that reflect the true evaluations of the users.

The Clarke Tax auctions are beneficial for decision making for multiple participants with different opinions, as they support truthfulness [32]. If the Clarke Tax auctions are applied in commerce, then each participant would have their own budget (e.g., money) with which to bid. However, in PANO, the participants are given points at the beginning of each auction, which they can use to bid in the current auction or save to bid later. As usual, a participant cannot bid more than his or her current budget, which is defined as the amount of unspent points. The PANO auctions have two constraints on bidding. The first is that, if an agent accrues points during an auction with a set of other agents, it can only use these points in a different auction with the same set of agents. This ensures that agents do not enter arbitrary auctions to accumulate points. The second is that each agent can bid in a specific range. This enables a common context for the given bids.

2.2 Privacy Personas

Online social networks are widely used throughout the world, with billions of users with varying understandings of privacy. These users differ in how they perceive privacy, which affects what they share online. Ideally, the personal assistants that are developed should be able to help users with different privacy preferences. In order to study this, it is beneficial to be able to categorize users into segments and to check if the personal assistants are beneficial for users in each segment. Westin conducted many surveys over decades and defined three categories of users in terms of privacy understanding, namely *Marginally Concerned*, *Fundamentalists,* and *Pragmatic Majority* [24]. Dupree et al. [13] extend Westin's categories according to their own qualitative study with surveys and interviews to define *privacy personas.* These privacy personas can be explained over two dimensions: *knowledge* and *motivation.* The knowledge dimension denotes how much a persona has knowledge about privacy choices in the system. For example, knowing whom a certain content can reach or knowing the implications of sharing a particular content would be ranked high in this dimension. The motivation dimension denotes how much effort a user is willing to expend to reflect his or her privacy choices in the system. For example, changing privacy

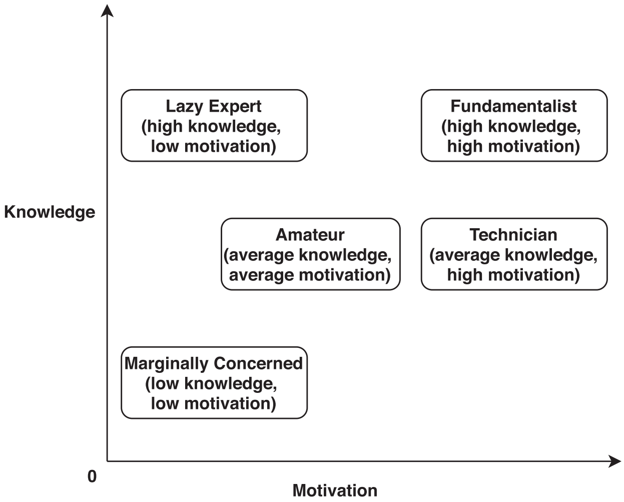
Knowledge
Lazy Expert (high knowledge, low motivation)
Fundamentalist (high knowledge, high motivation)
Amateur (average knowledge, average motivation)
Technician (average knowledge, high motivation)
Marginally Concerned (low knowledge, low motivation)
0
Motivation

Fig. 1. Knowledge and motivation dimensions for privacy personas, according to Dupree et al. [13].

settings for some content or spending time to check who can access content would rank high in this dimension. As depicted in Figure 1, Dupree et al. [13] organize five privacy personas over these two dimensions: *Fundamentalists*, *Lazy Experts*, *Technicians*, *Amateurs,* and *Marginally Concerned*. They place *Fundamentalists* and *Lazy Experts* higher on the knowledge dimension, while *Marginally Concerned* are the lowest. For the motivation dimension, again *Marginally Concerned* are the lowest, while the highest clusters differ to become *Fundamentalists* and *Technicians*. The categorization of Dupree et al. provides sufficient details for the classification of users, which helps us address how PANOLA can assist different types of users.

When personal assistants are making privacy decisions, they can consult their user to ask for *input* on whether a piece of content should be shared or not. Naturally, users who have more knowledge and motivation about privacy would give more input to their agents. Therefore, the agents of these users would have an advantage over others in the chosen privacy outcomes. It is known that the majority of a society consists of individuals with lower knowledge and motivation [13]. Ideally, the agents that we develop should help these users as much as those with high knowledge and motivation. Only then can the system establish *equity* and treat all users fairly [21]. To reach this, the software agents should be able to learn to bid in accordance with their users' privacy requirements over time, even when the input is sparse due to lack of motivation, or wrong when the knowledge of the users is not enough to provide input for correct decisions. To reach our equity goal, our learning agents should have three main capabilities. First, they should be able to learn from previous privacy decision outcomes to make better decisions in the future. Second, they should be aware of potentially wrong input by users, and not become confident about the privacy requirements of users with little input. Third, the agents should still be able to make decisions with little input, since some users might not have the motivation to even give input.

2.3 Running Example with Privacy Personas

Assume that Alice, Bob, Carol, Dave, and Emma are co-owners of a piece of content, which is going to be either shared or kept private based on the decision resulting from a PANO auction. For privacy personas, Alice is a *Lazy Expert*, who has extensive knowledge about privacy but lacks motivation to express her opinions for privacy decisions. Bob is classified as *Marginally*

Table 3. Example of a Decision by Various Privacy Personas

Name	Privacy Persona	Actual Preference	Expressed Preference
Alice	Lazy Expert *(low motivation, high knowledge)*	Not Share	-
Bob	Marginally Concerned *(low motivation, low knowledge)*	Not Share	Share
Carol	Fundamentalist *(high motivation, high knowledge)*	Share	Share
Dave	Technician *(high motivation, average knowledge)*	Not Share	Not Share
Emma	Amateur *(average motivation, average knowledge)*	Not Share	-

Concerned, which means that he has low-level motivation similar to Alice but also very limited knowledge about privacy. Carol is a *Fundamentalist*, which makes her highly motivated to express her privacy concerns over the system, and she also has high-level knowledge about privacy to be able to make appropriate decisions. Dave is a *Technician*; therefore, he has motivation similar to Carol but slightly less knowledge about privacy, which might cause him to make some mistakes while expressing his privacy concerns. Emma is an *Amateur* who has a similar knowledge level as Dave but is less motivated than him to express her privacy preferences.

Once the PANO auction commences, each co-owner would need to assess what privacy outcome would be more fitting to his or her privacy understanding and place a bid for that outcome in hopes of affecting the final decision in his or her favor. Table 3 shows an example setup with the given privacy personas for the participants above. Let's assume Alice, Bob, Dave, and Emma would want the content to be kept private, while Carol wants it to be shared. Since Alice is a *Lazy Expert* and lacks motivation, she does not spend time on the auction and does not place a bid, even though with her high knowledge she would have a clear idea of how to place a proper bid. Bob, being *Marginally Concerned*, also lacks motivation but still decides to participate. However, since he is not well versed in the auctions and the system, he places a small bid for sharing the content. Carol is a *Fundamentalist*; therefore, she is highly motivated to bid and also has the knowledge to back it up, and she places a bigger bid for sharing the content, knowing that there are three other co-owners, so forcing her decision might require her to bid a high amount. Dave is also highly motivated and has knowledge to an extent since he is a *Technician*; therefore, he places a bid for not sharing the content, but a bit less than Carol since he does not have the knowledge that none of the other agents might be bidding the same outcome as him. Emma, being an *Amateur*, does not express her preference due to not being strongly motivated according to her persona characteristics; hence, she does not bid for the auction, like Alice.

Next, the PANO auction is processed with the placed bids, and since share outcome bids outbid not sharing, the content is shared in the system, even though four of the agents would have preferred to keep the content private. Due to the accidental share bid of Bob, Carol is even taxed less because she was not the sole decision maker of the auction. This also would help her for future auctions, since she will still have some points to spend, and with her high motivation level, she would mostly be willing to spend time on the auctions for enforcing her privacy decisions over others.

Let's examine how PANOLA would be helpful to each user separately. Alice and Emma lack motivation and thus do not place a bid; but if they each had PANOLA, their personal assistants would have placed correct bids on their behalf. Bob lacks both the motivation and expertise. The PANOLA agent would learn to bid according to Bob's privacy expectations and place a correct

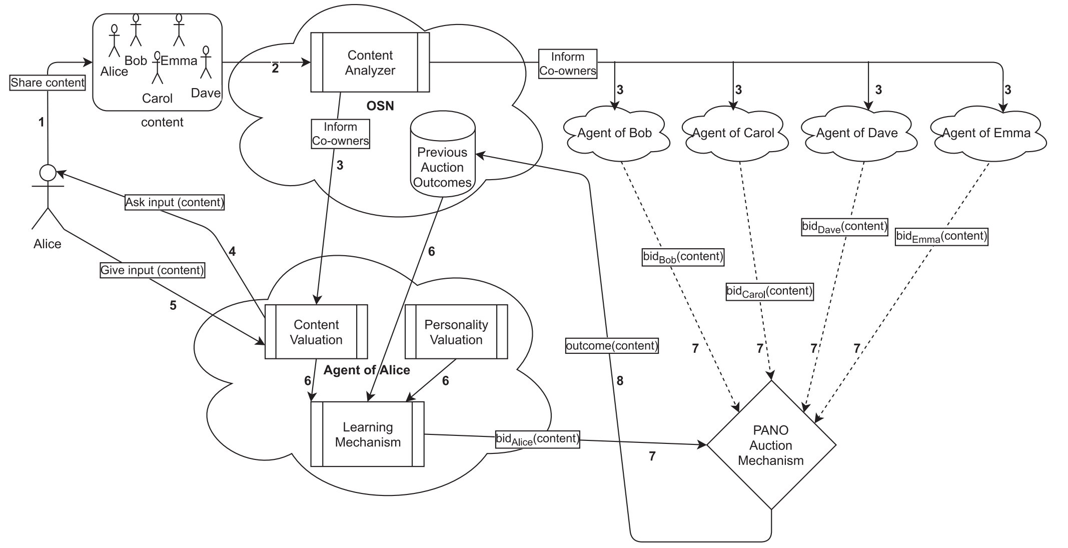

Fig. 2. Flow diagram depicting how a PANOLA agent learns and performs bidding outcomes.

bid accordingly. This would have avoided Bob's mistake. For Dave, the PANOLA agent would similarly estimate the correct bid and place it accordingly. Thus, the outcome of this auction would have been different, respecting the privacy expectations of users. This is expected to bring us close to enabling equity, where all the users are supported to enable each one to have an equal voice in collaborative privacy decisions while their privacy concerns are respected explicitly.

3 PRESERVING PRIVACY WITH PANOLA

In widely used OSNs, content sharing is done by a user who uploads the content onto the OSN, and it is shared either publicly or with a specific set of users, according to the user's choice. For a piece of content that is co-owned by multiple users, the remaining users can only have a say in the privacy decision after the OSN receives it to be shared. Instead, we advocate that an OSN provider should first identify the co-owners of the content (e.g., user tags), as well as the contextual properties of it (e.g., time of day or location), and then provide an opportunity for them to engage in decision making as to share or not share the content. Next, the personal assistants of all co-owners deliberate on the privacy outcome for the content in a distributed manner and act according to the outcome of the PANO auction.

Figure 2 depicts how we envision users and their personal assistants to act, according to both user input and the outcome of previous auctions. The numbers next to the arrows show the order of action, and the texts attached to the arrows depict the details of the related action. The rectangles with vertical lines on its right and left represent processes that receive input, deliberate on it, and produce an output, while a cylinder depicts data storage.

The flow starts with Alice wanting to share the image content on the top left (1), which is a group picture of her with Bob, Carol, Dave, and Emma. According to the diagram, when co-owned content requires a privacy decision, first, the content is analyzed to identify the co-owners and the contextual properties of it (2). Afterward, each agent of the co-owners is informed about these by the OSN provider (3). We omit the detailed actions taken by the agents of Bob, Carol, Dave, and Emma for brevity and only show this part of the flow for Alice, since all the agents go through the same process. After the personal assistants receive the information about the content, each assistant agent asks the users they represent for *user input* for the content decision (4), which would be just by informing the user about what type of content this is and requesting information about whether the user would like this content to be shared or kept private. Specifically, the agent interacts with the user to ask if a piece of content's preferred sharing outcome is *share* or *not share*. The user may or may not give input, and if given, the input may or may not be correct (5). These correspond to the "motivation" and "knowledge" dimensions of personas. In actual usage, the agent may wait for a predetermined time and then decide that the input is not given. Again, we show this process only for Alice in the figure, but since all the agents of the users have a similar process, we use dashed arrows for the other four for simplicity. As mentioned while explaining privacy personas, some users might not have the motivation to give input; then the agent can solely rely on the previous auction outcomes of the user (6) or the previous feedback received by the user for similar content. Making use of the available information, the agent decides on a bid according to what has been learned until that time, and it places this bid for the PANO auction (7). After all the co-owner agents place their bids, the PANO mechanism is triggered, and the outcome to *share* or *not share* the content is decided (8). The auction does not have to be synchronous and there could be a time window in which agents are expected to bid. In case a co-owner's agent is not available to bid (e.g., communication failure, late response), the PANO mechanism places zero bids for each outcome on behalf of the co-owner. After a privacy decision is reached with the PANO mechanism, the outcome is stored along with other previous auction outcomes, from which the agents learn to bid better with their internal learning mechanism. The details of the learning mechanism for the PANOLA agents will be explained in the following subsections.

3.1 Learning to Bid

An important aspect of PANOLA is to learn how to bid for a given user. Since users have different privacy preferences for different types of content, the bid that will be given for different sharing outcomes will vary. In addition, users might not have a clear understanding of privacy; therefore, inaccurate valuations might occur. Users of domains such as OSNs are usually not experts on privacy. Even though many users claim to be caring about privacy and think that they are able to express their privacy concerns, their actions can prove the opposite, which could even contradict their privacy understanding [1]. Thus, it is important to present privacy outcomes in a straightforward way. Following this, in this article, we consider two privacy outcomes: share and not share.

In order to facilitate a learning method where users can have a say in the outcome, we investigate various machine learning approaches. The first option would be to use a supervised machine learning approach. However, since every agent would require some feedback from an expert according to its own privacy understanding, it would be impractical in a highly dynamic privacy environment with a very large number of agents. Hence, using a supervised learning algorithm would be almost impossible. The second option is to use unsupervised learning. However, since every agent has its own decision mechanism with little input about their actions, it is also difficult to apply unsupervised learning methods to extract patterns or clusters. In PANO, agents do not know the privacy preferences of the other agents and only can see the resulting privacy outcome for an auction, their own bid, and the tax they pay afterward. Thus, they cannot obtain a clear view about the society and can only rely on the limited information they can access. Furthermore, large multiagent systems like OSNs or IoT environments can contain a high level of traffic for collaborative privacy decisions, where agents should decide on privacy outcome in a matter of seconds and with as little computation as possible. Agents usually work on limitations for hardware, broadband connection, and so on; thus, applying complex machine learning algorithms such as deep neural networks becomes unfeasible. The third option, which we adopt, is to use reinforcement learning,

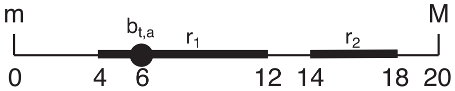

Fig. 3. A depiction of two ranges between minimum (m) and maximum (M) bidding boundaries and the initial bidding evaluation of content a for outcome t.

which enables agents to maximize their rewards from their actions [37]. When an auction is carried out, the outcome of the auction is used by agents as reward or punishment for the privacy action taken. Thus, an agent can model whether the taken action was useful for the given auction and if so reuse the same action later or switch to a different action if not.

3.2 Bidding Ranges

In a privacy auctioning mechanism such as PANO, picking the correct outcome and how much to bid for this selection is crucial. The agent needs to place a bid that reflects users' expectations in sharing. When learning how much to bid, the agent can aim to learn an exact bid value for a user or a range from which a bid will come. Learning an exact value is difficult because if the agent makes a mistake in bidding, it does not know if bids with close values would have sufficed. However, if the agent attempts to learn a range, then it can approximate the bids that express its user's preferences even if it cannot predict the bid precisely. Therefore, if an agent can learn ranges from which it can generate its bids over time, it would be easier to gradually get closer to the possible winning bids. For this reason, with the given minimum and maximum boundaries for PANO (m and M, respectively), we introduce bidding ranges, where the agents can pick from all the possible ranges within the boundaries and bid integers between the selected ranges.

Definition 3.1. A bidding range r is denoted as $[k, l]$ such that $k \geq m$, $l \leq M$, $k < l$, where m is the minimum and M is the maximum boundary for PANO. R denotes the set of all possible ranges $R = \{r_1, \ldots, r_n\}$.

Definition 3.2. For each auction outcome $t \in \{share, notshare\}$, each agent assigns a rational *utility* value to a range between 0 and 1 that denotes how beneficial a range r is for bidding for that privacy outcome, with 0 meaning the least suitable and 1 the most suitable (denoted as $Utility(r, t)$).

All the possible bidding ranges, outcomes, and their associated utility values are stored by the agents themselves in a suitable data structure to be maintained and updated as needed. When an agent is participating in an auction on content a, it generates a bid for the preferred outcome by first selecting the range that will yield the highest utility and then picking a value from this range. The bid given by an agent for content a for the outcome t is denoted as $b_{t,a}$. Picking $b_{t,a}$ from a range can be achieved according to a distribution function. We employ *Gaussian* distribution to pick $b_{t,a}$ from a selected range, which would favor the values that are closer to the mean of all integer values within the range. In some domains, a reduction in the number of ranges can be needed to reduce the computations. One solution could be to hold the PANO auctions with a small M so that the number of ranges decreases. Another solution could be to enforce a minimum length g on the range $[k, l]$ in Definition 3.1 such that $l - k > g$. Thus, for example, by having $g = 2$, we exclude some ranges, e.g., $[2 - 4]$ or $[7 - 8]$. Similarly, it could be possible to enforce a maximum length on the ranges to reduce imprecision. Here, we demonstrate our agent using all possible ranges.

Over time, utility values of bidding ranges change according to the success or failure of the selected bids. Agents do not share the utility values with the environment. Therefore, agents can update their utility values without letting the other agents know. Each agent updates its utilities according to the outcome of the auctions. Reinforcement learning is used to make agents learn to pick the most suitable range for a given content type, using information that results from the PANO auctions, such as the currency they paid according to their bids, the deducted tax amount if any tax was paid, and the outcome chosen by the auction, which can be considered as the most important factor for the learning process. We employ all these factors in our computations for learning the suitability of the ranges. The agents pick the range with the highest utility for a given content and bid an integer value inside this range according to their bidding strategy for their preferred outcome.

Example 1. Figure 3 depicts two bidding range examples ($r_1 = [4, 12]$ and $r_2 = [14, 18]$) for outcome t between minimum and maximum boundaries (m and M, respectively), assigned as 0 and 20. The actual set of ranges contains more than these two, since we include all possible integer ranges between m and M. $b_{t,a}$ shows the bid for outcome t, which is given as 6 (picked with *Gaussian* distribution) and means that the agent bid from r_1, if r_1 and r_2 were the only ranges for the agent, due to $Utility(r_1, t) > Utility(r_2, t)$.

3.3 Personalized Bidding

Utility values of the ranges change over time according to the outcome of the auctions. Agents pick the range having the highest utility value that they have sufficient budget to be able to bid from for the given content type. For simplicity, we explain our method over a single type of content, which means the context for different types of content is not considered. Like most of the traditional approaches in reinforcement learning [6, 11, 38], the unsuccessful range selections are penalized with a decrease in the utility value, while the successful ones have an increase in utility. In our approach, the utility of a range is based on the *effectiveness* of bids given from that range in previous auctions. Intuitively, an agent's bid has been effective if the outcome of the auction was the agent's preferred outcome while the agent did not bid too much and was not taxed too much. We formalize this intuition using two variables: value of content and value of conformism.

- Value of Content (V_{Ct}) captures how important a specific type of content is for a user. When a user considers some piece of content important, it means that the user wants its own privacy preference to be the final outcome for a collaborative decision at all costs. In this case, the agent of the user would be assertive about the amount of the placed bid. A lower bid might win the agent the auction; however, a higher bid would be less risky, since the others' bids are not known. The V_{Ct} is a factor when the agent is learning the minimum possible bid with which it can win the auction. A higher V_{Ct} would reduce the importance of the amount of the placed bid in the effectiveness calculation. In the opposite case, when some content is not important for a user, a lower value of V_{Ct} would enable the agent to fine-tune the placed bid in the learning process. With a lower V_{Ct}, the placed bid has more importance in the effectiveness calculation; therefore, the agent tries to learn the lowest possible winning bid. In this case, the agent might lose a few auctions while finding the winning bid, but since the type of content is not that important, it would help the agent to save the budget for future auctions where it might be needed for more important content.
- Value of Conformism (V_{Cf}) measures how much a user is willing to be similar to the rest of the population. In the PANO auctions, if a participant places a bid that does not change the outcome, then we consider this as conforming to the group. Implicitly, V_{Cf} determines the extent of the tax an agent is willing to pay for an auction. Paying a tax means that the participant made a decision different from some other participants. With V_{Cf}, the effect of the paid tax comes into consideration for effectiveness calculation. With a higher V_{Cf}, agents would value the paid tax in the learning process, trying to minimize it. On the contrary, when

Table 4. Values for Utility Calculations

Name *Abbreviation*	Short Description	Equation/Function	Range
Value of Content V_{Ct}	Used for distinguishing between winning with lower and higher bids	$V_{Ct} \rightarrow 0$: increase effect of V_{Ct} $V_{Ct} \rightarrow 0.5$: decrease effect of V_{Ct}	[0,0.5]
Value of Conformism V_{Cf}	Changes the importance of taxes in utility calculation	$V_{Cf} \rightarrow 0$: decrease effect of V_{Cf} $V_{Cf} \rightarrow 0.5$: increase effect of V_{Cf}	[0,0.5]
Confidence C_t	Used for defining how confident the agent is about user input for outcome t	$C_t = C \times c_t/p$ $(C = 1 - e^{-p/S})$	[0,1]
Effectiveness E	Calculates effectiveness of a range	$E(r) = 1 - \left((0.5 - V_{Ct}) \times \frac{b_{t,a}}{M} + V_{Cf} \times \frac{Tax}{M}\right)$	[0,1]

an agent has a lower V_{Cf}, it will act to have its decision to be the final one and pay tax for it accordingly.

Table 4 summarizes the important parameters for the proposed approach, namely the *value of content*, *value of conformism*, *confidence*, and *effectiveness*. Recall that the user input received by the agent does not always reflect the user's actual privacy preferences, especially for users that lack knowledge on privacy. To address this, we introduce a confidence value that enables the agent to evaluate how confident it is about the user input. The parameters to compute confidence value C_t for the privacy outcome t are c_t, which is the count of user input that has been received in favor of outcome t, p, which holds how many times an input is received from the user, and S, which is the stability value that denotes the number of received user input to consider for making confident decisions. To compute the confidence value C_t, we first calculate a confidence coefficient C according to Equation (1):

$$C = 1 - e^{-p/S} \tag{1}$$

We adopt Equation (1) as a variation of the aging curve formula from the literature [44]. Our equation differs in the way that while the original aging curve value starts from 1 and decreases over time, the confidence value in our equation starts from 0 and increases over time. According to Equation (1), C value starts from 0, when the total number of user input received (p) is zero. Then, it will start to increase from 0 to 1 with every incoming input from the user. The confidence value will get closer to 1 after the stability value S is reached with the number of input (p). S should be set according to domain requirements, where an agent with a high S value will require a high number of user input to establish its certainty. Before the stability value is reached, the value change for C is steep, while the changes become slower after that point. After experimenting with several S values, we have assigned S for all our experiments as 10. The reason for this decision is that with a lower S value, the agents prematurely become confident about user input, which in some cases results in learning wrong outcomes, especially when the user lacks knowledge about privacy decisions. In the opposite case, higher S values slow down the increase of confidence, which causes the agents to bid less, while decision changes rarely occur after p is higher than 10. With the confidence coefficient equation in place, the confidence value C_t can simply be calculated with the equation below, which is achieved by multiplying C with the ratio of the number of input in favor of outcome t to the total number of input by the user (p).

$$C_t = C \times c_t/p \tag{2}$$

When the outcome of an action is the same as the outcome the agent bid for, the effectiveness E of this bid $b_{t,a}$ chosen from a range r for outcome t depends on the amount of the placed bid, the amount of tax received for content, and conformism valuations (V_{Ct} and V_{Cf}). After the auction, all these values are known by the agent, and the effectiveness can be calculated with Equation (3):

$$E(r) = 1 - ((0.5 - V_{Ct}) \times b_{t,a}/M + V_{Cf} \times Tax/M) \tag{3}$$

For the Effectiveness (E) value, a higher amount means that the agent's preferred outcome has been chosen with a lower bid and low tax. The ratio of $b_{t,a}$ to the maximum possible bid M gives the magnitude of the bid. The higher this value, the less effective the auction will be. This magnitude is adjusted with V_{Ct} to account for the fact that different agents would care about this differently. The ratio of Tax to maximum possible bid M gives the magnitude of the budget loss for the agent. Again, the higher this amount, the less effective the auction would be. Adjusting it with V_{Cf} enables the agent to account for different contexts, e.g., when the agent values some content a lot and would not want to conform with the others.

The effectiveness of a range will determine the likelihood of a bidding range to be selected again. With the V_{Ct} and V_{Cf} values, we ensure that agents can adjust their learning strategy according to the importance of the content and their will to conform with others. The highest possible bid would be the optimal strategy for a one-shot auction, since the leftover budget would not have any use, leaving the only goal as winning the auction. However, it would be costly in recurring PANO auctions, since it might cause the agent to pay a significant amount of tax along with a high bid when its bid has an impact on the auction outcome. But if the agents try to minimize the amount of bid and the possible tax for the winning bid, setting V_{Ct} and V_{Cf} accordingly can help them to save budget for future auctions.

3.4 Utility Update

When updating the utility of a range, there are two important sources of information. The first is what the agent has learned about the range based on the effectiveness calculations (Equation (3)) from previous bids. The second is its user's input on the preferred outcome. However, since some users are not knowledgeable in privacy, as depicted in the personas, the agent needs to model the confidence it has in its user for different outcomes (C_t), which can be obtained with the confidence calculations (Equation (2)).

The utility of a range r for a preference outcome t is then computed with the formula below:

$$Utility(r,t) = \left(\frac{\sum_{i=1}^{n} E_i(r)}{n} + (1 - |C_t - \hat{mean}(r)|) \right) /2 \tag{4}$$

According to Equation (4), for any range r, the utility value is the average of two values: the effectiveness average of all previous n number of privacy bids made within the range and the distance of confidence value gained after the feedback from the user (C_t) to the normalized value of the mean of the range values ($\hat{mean}(r)$) over the bidding boundaries. The distance calculation for the right-hand side of the formula ensures that when the confidence value is high, the agent would prefer to make bids from ranges with values closer to the maximum boundary M, since it would be more confident about the privacy preferences of the user. On the opposite case where the confidence is still low, the agent would prefer to bid from ranges with values closer to the minimum boundary m, since a higher bid would be risky because the agent would not be sure that its choices are in line with the user. The utility update ensures that both results of the previous auction outcomes and the user input are considered. We give an example below to demonstrate how the confidence value would affect the selected range.

Table 5. Utility Update Examples for Example 2

Previous Input	Confidence	Effectiveness	Distance of Range from Confidence	Utility
1 *not share*	$C_{notShare} = 0.1$	$r_1 = 0.3$ $r_2 = 0.9$	$r_1 : 1 - \|0.1 - 0.15\| = 0.95$ $r_2 : 1 - \|0.1 - 0.85\| = 0.25$	$\boldsymbol{Utility(r_1, notShare) = (0.3 + 0.95)/2 = 0.625}$ $Utility(r_2, notShare) = (0.9 + 0.25)/2 = 0.575$
8 *not share,* 2 *share*	$C_{notShare} = 0.5$	$r_1 = 0.3$ $r_2 = 0.9$	$r_1 : 1 - \|0.5 - 0.15\| = 0.65$ $r_2 : 1 - \|0.5 - 0.85\| = 0.65$	$Utility(r_1, notShare) = (0.3 + 0.65)/2 = 0.475$ $\boldsymbol{Utility(r_2, notShare) = (0.9 + 0.65)/2 = 0.775}$

Example 2. Let us consider two cases. In the first case Alice gives a single input for not sharing, while in the second she gives input for 10 times, 8 for not sharing and 2 for sharing. The summary of the utility calculations of the two cases in this example can be seen in Table 5. Using Equation (2) and stability value as $S = 10$, the confidence value of not sharing action (C_t where t is *notShare*) for the first case can be computed as $C_{notShare} = 0.1 \times 1/1$, equaling to 0.1. For the second case, the same value would be calculated as $C_{notShare} = 0.63 \times 8/10$, which would yield 0.5. Let's also assume that the boundaries (m for the minimum boundary and M for the maximum boundary) for bids were $m = 0$ and $M = 20$. Alice previously tried only two ranges for the auctions, $[0, 6]$ and $[14, 20]$, which are represented as r_1 and r_2, respectively. For both cases with varying input, the effectiveness average of the first range was 0.3 and the latter was 0.9. When the range means are normalized, r_1 would equal to $\hat{mean}_{r_1} = 0.15$ $(3/20)$ and $\hat{mean}_{r_2}$ would be 0.85 $(17/20)$. The distance from the confidence value for the first case then will be calculated for r_1 range as $1 - |0.1 - 0.15| = 0.95$, meaning that confidence is highly matching with this range. However, r_2 for the first case would be $1 - |0.1 - 0.85| = 0.25$, which means the agent is still not confident enough to bid that high. For the second case, recall that the confidence value was 0.5 instead of the 0.1 computed for the first case with less input. Thus, for this case the same calculations would give $1 - |0.5 - 0.15| = 0.65$ and $1 - |0.5 - 0.85| = 0.65$, which indicates the agent is equally confident for bidding both ranges. Since the final utility would be computed with the mean value of the confidence valuations and the effectiveness values, the first case would have $Utility(r_1, notShare) = (0.3 + 0.95)/2 = 0.625$ for range r_1 and $Utility(r_2, notShare) = (0.9 + 0.25)/2 = 0.575$ for range r_2. Therefore, with the bigger utility for range r_1, the first case would result in the agent preferring this range over r_2. For the second case, utility values would be calculated as $Utility(r_1, notShare) = (0.3 + 0.65)/2 = 0.475$ for range r_1 and $Utility(r_2, notShare) = (0.9 + 0.65)/2 = 0.775$ for range r_2. Hence, in the second case, since the agent is more confident because of receiving more input from Alice, the r_2 range would be favored for the auction bid.

3.5 Decision Making with PANOLA

PANOLA agents employ the utility formula explained in the previous subsection to make decisions on which bidding range should be chosen for the current PANO auction. As shown in Figure 2, agents interact with the OSN to make privacy decisions, since the PANO auctions are governed by the OSN.

Algorithm 1 explains the steps the PANOLA agents take throughout the decision-making process. When a new piece of content is introduced to the OSN that requires a collaborative decision, agents are informed about it. In line 1, the agent asks its own user about an input of *share* or *not share* about content a. If an input about the outcome is received from the user, the confidence values for that outcome are updated (Equation (2)), as seen in lines 2 to 4. Lines 5 and 6 set an initial range and outcome for the agent, which are r_0 (an arbitrary $r \in R$) and *not share,* respectively. Then, for each possible range, utilities are updated according to Equation (4) (line 8) as explained in Section 3.4. This update operation is necessary because a possible change in the confidence or effectiveness values will yield a new utility value for the same range and outcome. Then, the range with the related outcome that has the highest utility value is selected for bidding (lines $7 - 13$). In line 14 an integer bid is chosen from the selected range, and this bid is sent to the PANO mechanism

ALGORITHM 1: PANOLA agents

```
   Parameter : $R$: Set of ranges
   Parameter : $T$: Set of outcomes $\{share, notShare\}$; outcome $t \in T$
   Input: $a$: Content
   Data: $C_t$: Confidence for outcome t
   Data: $E$: Set of effectiveness values for each $r \in R$
   Output: $b_{t,a}$: Bid for outcome $t$
1  ask(input($a$), user)
2  if (input(a) exists) then
3      update($C_t$,input($a$)) //Equation (2)
4  end
5  $best \leftarrow r_0$
6  $outcome \leftarrow notShare$
7  foreach $r \in R$ do
8      update(Utility($r, t$)) //Equation (4)
9      if Utility($r, t$) $\geq best$ then
10         $best \leftarrow r$
11         $outcome \leftarrow t$
12     end
13 end
14 $b_{t,a} \leftarrow$ bid($best$,$outcome$)
15 send($b_{t,a}$,PANO)
16 receive($< decision, tax >$)
17 $E \leftarrow E \cup E(best)$ //Equation (3)
```

in line 15. In some cases, due to agents accruing bidding points separately for every different set of co-owners, the agent might not have enough budget to bid for the range with the highest utility for a given set of co-owners. In this case, the selected range becomes the highest possible one that the agent is able to bid from with its owned currency. Afterward, the PANO auction commences when all the bids are in place from all the co-owners. The resulting decision of the PANO auction is received by the agent along with the amount of tax to be paid (line 16), and the agent adds the effectiveness of its recent bid ($E_{(best)}$) (Equation (3)) to the set of all previous effectiveness values (E) in line 17.

3.6 Running Example with PANOLA

We now walk through our running example from Section 2.3, but now with the PANOLA agents employed to represent the users. Since PANOLA agents learn to bid over time according to input from the users, the decisions might differ based on the number of previous privacy decisions made by the same co-owners. To represent this, we show two executions, first the initial auction where none of the users joined an auction together before, and one after 20 auctions together. We assign the stability value for the confidence as $S = 10$, and the range boundaries as $m = 0$ and $M = 20$, while the possible outcomes are *share* and *not share*. Table 6 and Table 7 show the information about the users and their personas, their prior input, and for which output they placed their bids for both examples, respectively.

In the first iteration of the example in Section 2.3, the agents have no prior information. Bob, Carol, and Dave are the users to give input as seen in Table 6. Recall that Bob has the *Marginally*

Table 6. Example of a Decision with No Prior Knowledge

Name	Privacy Persona	Input	Actual Preference	Confidence	Bid
Alice	Lazy Expert	-	Not Share	0	-
Bob	Marginally Concerned	Share	Not Share	$C_t = 0.095$	$2, t = Share$
Carol	Fundamentalist	Share	Share	$C_t = 0.095$	$2, t = Share$
Dave	Technician	Not Share	Not Share	$C_t = 0.095$	$2, t = NotShare$
Emma	Amateur	-	Not Share	0	-

Table 7. Example of a Decision after 20 Previous Auctions

Name	Privacy Persona	Input	Actual Preference	Confidence	Bid
Alice	Lazy Expert	0 Share 7 Not Share	Not Share	$C_t = 0.50$	$10, t = NotShare$
Bob	Marginally Concerned	2 Share 6 Not Share	Not Share	$C_t = 0.41$	$7, t = NotShare$
Carol	Fundamentalist	18 Share 0 Not Share	Share	$C_t = 0.83$	$16, t = Share$
Dave	Technician	3 Share 14 Not Share	Not Share	$C_t = 0.67$	$13, t = NotShare$
Emma	Amateur	3 Share 10 Not Share	Not Share	$C_t = 0.56$	$11, t = NotShare$

Concerned persona and not sufficient knowledge about privacy; he accidentally gives input to his agent to share it while he actually would have wanted it not to be shared. Carol is the *Fundamentalist* and knowingly advises the same to her agent. Dave's input to his agent is for not sharing the content. In this case, Alice's agent cannot place a bid since Alice is a *Lazy Expert* with low motivation to participate in decisions, causing the agent to have no prior input. Similarly, Emma does not give input, even though she is slightly more motivated than Alice, yet not as much as Carol or Dave. Even though the agents that represent Alice and Emma can take initiative to bid without input from them, a randomized bid would be more harmful to their privacy than not bidding since there is a possibility that the agents can bid the opposing outcome without prior input. Therefore, for both Alice and Emma, leaving the decision to other participants would be the preferred choice rather than taking the risk of bidding against their actual preferred outcome. If we assume that all the agents have the same value settings, the bidding agents would have the same low confidence ($C_t = 0.095$, if the stability value S is assigned as 10) for the user input as this is the first auction. Therefore, since this confidence value would result in a selection from lower bid ranges (i.e., the ranges that contain lower bid values), Bob and Carol's agents bid a low amount (i.e., a bid closer to the minimum boundary for the ranges) for sharing the content, while Dave places a similar amount to keep the content private. In this case, the outcome would still be the same as in the case where the users would bid for themselves, due to the wrongly placed bid of Bob.

With more input from the users over time, the agents gain confidence concerning their users' privacy requirements and therefore can bid better on their behalf. According to Table 7, after 20 auctions with the same co-owners, we assume that Carol and Dave are the ones who give the highest number of input (18 and 17 input, respectively, for this example), so their agents would be more confident ($C_t = 0.83$ and $C_t = 0.67$, respectively, if the stability value S is assigned as 10) to bid higher values for the presumed privacy action. Emma would not reach the same confidence value due to having a lower motivation than Carol and Dave, but it would still have an average level of confidence ($C_t = 0.56$, with the stability value S assigned as 10) since her levels of motivation and knowledge are considered average. With the lower motivation of Alice and Bob, the number of inputs received from them are fewer than that of Carol, Dave, and Emma (7 and 8 input, respectively), but wrongly placed input like Bob's first one is filtered over time, since even with less knowledge, every persona type would still tend to give input according to their actual privacy understanding more than the wrong input. According to our confidence formula, both Alice and Bob would have lower confidence values due to less number of inputs and some wrong input by Bob ($C_t = 0.50$ and $C_t = 0.41$, respectively, when the stability value S is assigned as 10), as can be seen in the confidence column of Table 7. Thus, their agents would bid a lower amount for the preferred privacy action of their users. When the auction is commenced, Alice, Bob, Dave, and Emma would bid for not sharing the content, and Carol would bid for sharing. With these bids, the outcome will be to *not share* the content. We consider this to be a good decision as four out of five users' privacy requirements are satisfied and these four individuals influence the outcome to the extent that they care about the decision through their bids. The best outcome is when everyone's privacy preferences are satisfied. However, in reality, when multiple co-owners exist, this is rarely going to be the case. Note that with majority voting the result would have been the same, but PANO advocates that individuals for which the content is worth more should influence the outcome more. However, those agents are left with less to spend in the following auctions. Therefore, when all PANOLA agents learn to bid as in this example, the privacy of everyone will be preserved in the long run.

4 EVALUATION

In an OSN, users with varying privacy understanding would be classified under different privacy personas, and PANOLA agents who represent these users aim to represent them regardless of their differing knowledge and motivation levels. Over time, each PANOLA agent learns from previous collaborative decisions and the input received from its user. After a sufficient number of input and prior knowledge, PANOLA agents become confident about the privacy understanding of the users they represent and bid in a manner to reach a decision in their represented user's favor, while not overbidding so that they would still be able to have a say in future privacy decisions. Each agent might require a different number of decisions to reach that level, since some users might not have the motivation to provide input, or they might not have enough knowledge to express their privacy requirements correctly. In spite of these learning differences, after a certain number of collaborative decisions, each agent should have learned their users' privacy requirements, and afterward an equity of the decisions should be seen.

Research Questions: We formulate the following research questions:

- **RQ-1**: Can PANOLA agents learn to bid correctly (i.e., a possible winning bid for the preferred outcome of the user), thus improving how well users preserve their privacy?
- **RQ-2**: Do PANOLA agents help users of different types (e.g., those who know less about privacy than others) well, thereby leading to equity of treatment?
- **RQ-3**: Can PANOLA agents help others preserve their privacy by finding the right balance between individualism and conformism?

Simulation Setup: We use multiagent simulations to study these questions. Each PANOLA agent in the multiagent simulation represents OSN users with various privacy personas and bids on behalf of them. The setup for the simulation is as follows: First, a number of users with varying privacy personas are introduced into the simulation. The personas employed are in line with the Dupree et al. [13] five persona types, which differ in knowledge and motivation dimensions. The number of users belonging to one of these personas is determined probabilistically, again in line with the percentages of personas found in the Dupree et al. studies, such that the probabilities for a user belonging to a persona are 23% for *Marginally Concerned*, 34% for *Amateurs*, 18% for *Technicians*, 21% for *Lazy Experts,* and 4% for *Fundamentalists*. Even though these percentages are expected to represent the privacy personas of real-life OSN users, they could differ depending on the application domain or other factors. We would still expect our research questions to have similar results as our research questions target individual agents, rather than the interactions between different types of agents. Along with privacy personas, we also employ a *random agent* for some tests to showcase what the baseline performance would be. The *random agent* bids a random integer amount within the boundaries for either *share* or *not share* outcome, which is again chosen randomly, for each auction in which it participates. The *random agent* does not make use of the user input; therefore, the placed bid can either be in favor or against the user's privacy requirements.

We define knowledge and motivation on a scale from 0 to 100, with 0 representing the lowest knowledge and motivation and 100 the highest. To be in line with knowledge and motivation dimensions represented in Dupree et al. [13], we assign three levels for both, which can differ for each persona. Table 8 shows the knowledge and motivation levels for each persona, as well as their percentage in the entire community. According to Table 8, the *Marginally Concerned* have both the lowest knowledge and motivation levels; therefore, we assign a value of 10 for both. With this value, the *Marginally Concerned* are motivated to give feedback for only 10% of the privacy decisions, and only 10% of these decisions are given correctly while the rest are random, since they have a low level of knowledge. For *Amateurs*, motivation and knowledge values are assigned as 40, since they are defined as higher in both dimensions than *Marginally Concerned*. *Technicians* again have 40 set as the knowledge value, but their motivation is higher than *Amateurs*; thus, we assign it as 70. *Lazy Experts* lack motivation as much as *Marginally Concerned*; thus, the value is again set for 10. However, they have higher knowledge than other personas except *Fundamentalists*, which we assign as 70. *Fundamentalists* share the same level of knowledge with *Lazy Experts*, but they also have much higher motivation; therefore, we assign the value of 70 for both dimensions for *Fundamentalists*. The stability value S is assigned as 10 for the confidence calculations. The V_{Ct} and V_{Cf} values are both assigned as 0.5 by default for all agents, which supports a balanced strategy between individualism and conformism, since the agents would both try to minimize their bids and their taxes. After personas and their aforementioned knowledge and motivation levels are set, we introduce co-owned content to the simulation, which requires a collaborative decision with the PANO mechanism. From the set of users, co-owners are assigned to these pieces of content. Each piece of content has a random number of co-owners, differing from 2 to 5. In order to reduce sparsity and have more interactions with similar co-owners, the population for possible co-owners is set to 20. Then, for each user, a PANOLA agent is assigned. These PANOLA agents ask for input from the users when a privacy decision is going to be made for some content and might receive input depending on the levels of motivation by users. The received input can also differ from the actual preferred privacy outcome of the users, since all users have varying knowledge levels. In the light of this setup, content decisions are made sequentially. The simulation environment is developed in Java, and the simulations are run with Eclipse IDE 4.14 on Windows 10 OS and with an Intel i7-6700HQ processor. The settings for the experiments can be obtained from https://git.science.uu.nl/o.ulusoy/panolasim.

Simulation Metrics: During the simulations, PANOLA agents accumulate user input over time and learn from the previous PANO-based collaborative privacy decisions. Our main metric for evaluation is success, which measures how successful an agent is in the actions in which it participates. We consider an action taken by an agent as successful when the agent bids according to the actual privacy outcome the user wants, and the outcome of the auction is in favor of the user. Note that

Table 8. Percentage of All Personas in Our Evaluations and Their Knowledge and Motivation Levels

Privacy Persona	Percentage in Community	Knowledge Level	Motivation Level
Marginally Concerned	23%	10%	10%
Amateur	34%	40%	40%
Technician	18%	40%	70%
Lazy Expert	21%	70%	10%
Fundamentalist	4%	70%	70%

a user might not always give input that would lead the agent to the correct outcome, and in that case the agent might bid for an outcome that is not the actual intent of the user. In that case, even if the agent wins the auction, we consider the outcome unsuccessful. We denote the total number of successful actions by an agent as SA and the number of auctions the agent participated in as n to define the success metric s, which is simply calculated with the formula below:

$$s = \frac{SA}{n} \tag{5}$$

We also investigate the statistical significance of the results by computing confidence intervals of the final results for each setup. A confidence interval is a range of values calculated by statistical methods that include the desired true parameter (the arithmetic mean in our case) with a confidence level. Confidence intervals suit our experiment well [12], since they demonstrate the probability of the deviation from the results. In order to evaluate the confidence intervals, we employ the following formula:

$$\overline{X} \pm Z \times \frac{\sigma}{\sqrt{n}} \tag{6}$$

where $\overline{X}$ is the mean of the results, Z is the value obtained from the z-score for the selected confidence interval, σ is the standard deviation of the results, and n is the number of runs. For the confidence intervals, we select 95% as our confidence value as it is common practice in many scientific experiments [12]; hence, the Z value is assigned as 1.960.

4.1 PANOLA Against Non-Learning Agents

With the simulation setup explained above, we first investigate whether PANOLA agents can learn to bid correctly, i.e., place a bid that would yield a result in favor of the user's privacy preference against non-learning agents that employ a fixed bidding strategy that do not involve learning over time. We name this strategy as ***simple bidding scheme*** (*SBS*), with which the agents would bid a predetermined value of 10 if they have sufficient budget and, if not, their current budget. *SBS* constitutes the base case to test our PANOLA agents. In our setup, we set the PANO bid boundaries as $[0, 20]$ and the amount of earned budget for an auction as 10. PANOLA agents act as explained in Section 3.1.

We run two experiments to evaluate the success of each privacy persona represented by PANOLA agents or non-learning agents. For each setup, we execute 50 runs for each persona, each with 11,000 co-owned content decisions. To determine the number of runs, we conducted a preparation experiment with some of the setups we evaluate in our experiments, where we compared the results for a given scenario with 25, 50, 75, and 100 runs, respectively. The outcome of this experiment showed that the deviation in the results becomes less than 0.5% when the number of runs is 50 or higher. Since we also will demonstrate confidence intervals with our results, we have decided to adopt 50 runs for each setup. Other co-owners are assigned from 20 agents in the network, which can be one of the five personas according to the values in Table 8. In addition to five personas, we run the same setup with a *random agent* to compare the performance

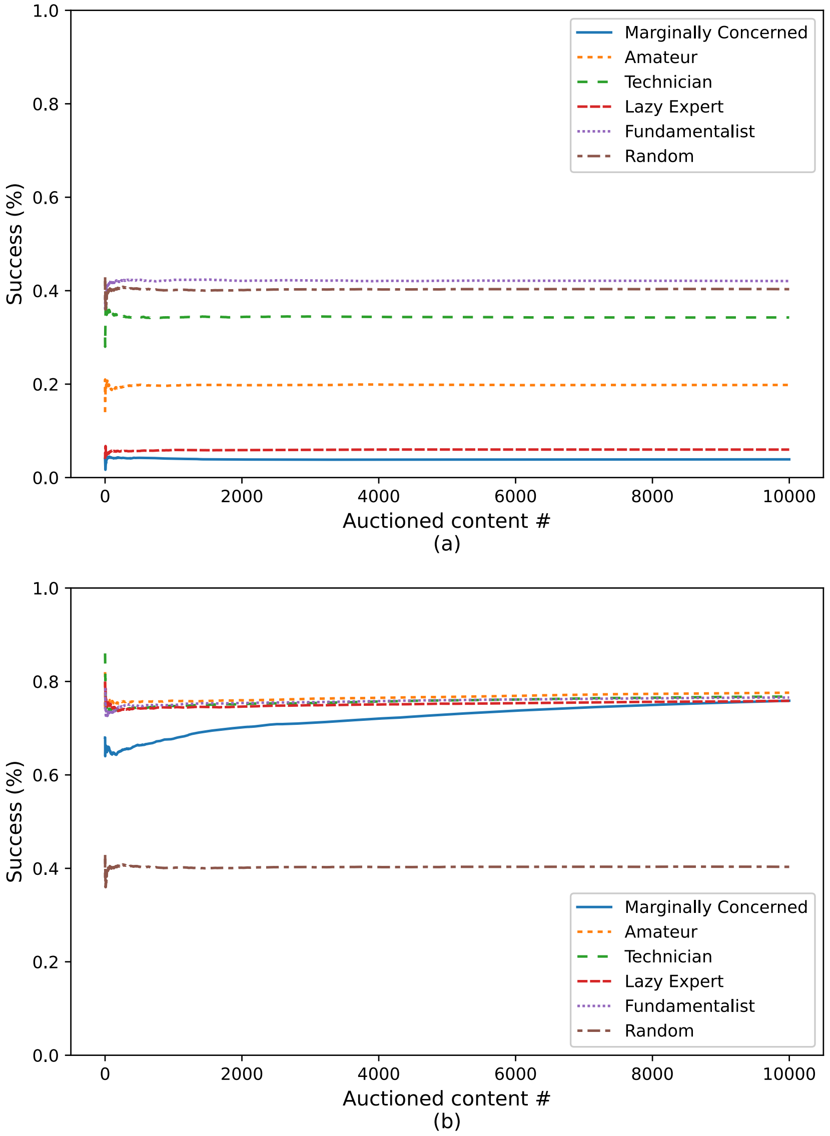

Fig. 4. Success of (a) non-learning agents against each other and (b) PANOLA agents and *random agent* against non-learning agents.

of PANOLA agents with this base case. We use the first 1,000 items of content for the learning phase of PANOLA agents and therefore do not measure the success for those. The auctions for the remaining 10,000 items of content are used for the test set; thus, they are a part of the evaluations.

Figure 4 shows the results of experiments with the given setup. The graph on the left (a) shows the success of each privacy persona type and the *random agent*, when all the agents in the simulation are non-learning and employ *SBS*. The graph on the right (b) shows the success of PANOLA agents who learn privacy ranges over time and the *random agent* engaging into auctions with non-learning agents. The success values on the figures are presented after the first auction until the outcome of the last auction. Therefore, the lines that depict the percentage of success do not include any information for 0 on the x axis.

According to the results in Figure 4(a), when agents do not learn from input or previous privacy decisions, the success of all personas except the *Fundamentalists* is below the success of *random agent*, meaning that these personas lose most of the auctions they enter and perform worse than even a random bidding strategy. This is mainly caused by *random agent* bidding for every auction while the agents can only bid when there is an input from their users, which happens only occasionally according to the motivation levels and can still be wrong especially for lower knowledge levels. The *Fundamentalists* perform slightly better than the *random agent* as they have both the knowledge and the motivation. This result demonstrates that there is indeed a need for a personal assistant that can learn from the user and help them preserve their privacy for the other personas. Otherwise, users who have more privacy knowledge or higher motivation to express their privacy preferences can dominate the privacy decisions in their favor, making it impossible to reach *equity*. Note that due to the mechanism in PANO, reaching higher percentages of success is extremely difficult. This is because if a participant's bid is affecting the final decision, he or she also gets taxed, which would leave him or her short-handed for the next auction since there might be no points to spend for another auction. In relation to this, *random agents* would not necessarily reach 50% success either, as we observe in this experiment where the *random agent* can only reach 40% against non-learning agents and performs even worse against learning agents.

The results in Figure 4(b) show that all personas perform successfully when PANOLA agents learn to bid over time, where each type has a success percentage above 75% after 11,000 auctions. Since the agents already go through a learning phase for the first 1,000 items of content, their success quickly reaches a stable point, with the exception of the *Marginally Concerned*, which still performs fairly well with a success above 65% at start but reaches the performance of other personas after processing 6,000 items of content. This is due to the *Marginally Concerned* rarely giving input and these inputs often being wrong due to the persona's low knowledge level. We also observe that a *random agent* becomes inferior to PANOLA agents. Referring back to RQ-1, we can say that PANOLA agents learn to bid correctly over time and thus help users preserve their privacy.

Comparing the non-learning and learning agents in Figure 4, we also observe that PANOLA agents greatly improve the success of all persona types, not only of those with higher motivation or knowledge. In the non-learning setup, there are big gaps between the success of various personas, which means that some personas have less say in the privacy decisions than the others. When PANOLA agents are employed, regardless of the privacy persona of the users, each agent can reach a similar success percentage in the end, which means that each of the users represented by PANOLA agents won a similar number of auctions; hence, each user was treated equally regardless of their personas with varying knowledge and motivation levels. Therefore, we answer RQ-2 positively, such that PANOLA agents help users of different types (e.g., those who know less about privacy than others) well, thereby leading to equity of treatment.

Table 9 shows our results of statistical significance with confidence intervals for the final results of this setup. According to Table 9, we can see that the experiments from Figure 4, where non-learning or PANOLA agents are evaluated against non-learning agents, all have confidence intervals less than 2%, which shows that the expected outcome of each run would be similar to our presented results.

4.2 PANOLA Agents Against Each Other

We have shown that PANOLA agents can learn to bid better over time, but can they still manage to bid correctly when the other co-owners also employ PANOLA agents for auctions? To test this,

Table 9. Statistical Significance Analysis of Experiments Depicted in Figure 4, after 50 Runs per Experiment

Agent Type	Learning Type	Mean of Success	Confidence Interval (95%)
Marginally Concerned	Non-learning	3.89%	±0.09
Amateur	Non-learning	19.82%	±0.39
Technician	Non-learning	34.26%	±0.68
Lazy Expert	Non-learning	5.99%	±0.13
Fundamentalist	Non-learning	42.06%	±0.87
Random	Non-learning	40.32%	±0.56
Marginally Concerned	Learning	75.87%	±1.75
Amateur	Learning	77.58%	±1.04
Technician	Learning	76.76%	±1.10
Lazy Expert	Learning	75.84%	±1.19
Fundamentalist	Learning	76.56%	±1.14
Random	Non-learning	40.32%	±0.56

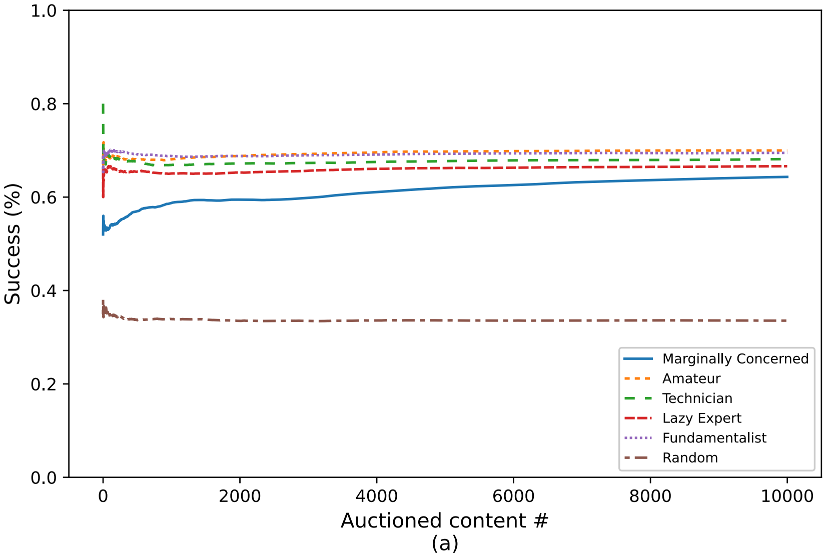

(a)

Fig. 5. Success of PANOLA agents against other PANOLA agents with varying privacy personas.

we introduce another experiment, similar to the one in the previous section, but only including PANOLA agents as co-owners, which are again assigned between 2 to 5 for each piece of content. Again, we measure the success of personas separately. Therefore, the agent of which we measure the success has a predetermined persona, while the others can be one of all five personas, according to the probabilistic values given at the beginning of this section. We additionally test the *random agent* for comparison against PANOLA agents. Again, we have 11,000 items of content, 1,000 for training and 10,000 for test of success, and we perform 50 runs with the same setup for each persona.

Figure 5 shows the results of this experiment. The trend of success is similar to the results from the PANOLA agents against non-learning agents setup, but each persona has slightly lower success percentages due to other agents also being PANOLA agents as they also learn to bid over time. However, agents that represent each persona manage to be successful in more than 60% of

Table 10. Statistical Significance Analysis of Experiments Depicted in Figure 5, after 50 Runs per Experiment

Agent Type	Mean of Success	Confidence Interval (95%)
Marginally Concerned	64.32%	±3.80
Amateur	69.98%	±3.11
Technician	68.14%	±3.46
Lazy Expert	66.61%	±3.44
Fundamentalist	69.47%	±3.39
Random	33.54%	±1.52

the auctions, which is significantly better than the case where all agents are non-learning. This shows that, even when the other agents are adaptive in their bids, PANOLA agents can also adapt over time to keep placing winning bids for every privacy persona with varying knowledge and motivation levels. Again, the non-learning agent that bids randomly performs the worst, with a success rate of only about 35%. We can still see that success percentages of all personas are close to each other after 10,000 auctions, while the *Marginally Concerned* are reaching that point slightly slower than the other personas due to low knowledge and motivation levels to provide sufficient information to the agents. Moreover, we can still say that the results are in line with our goal of equity, since every privacy persona has a similar number of successful outcomes to have a say in the collaborative privacy decisions.

Table 10 depicts the results of statistical significance analysis with confidence intervals for this experiment. When all the agents are PANOLA agents, the margin of error slightly increases over the results of the previous subsection. However, the confidence intervals for each persona are still less than 4%, which shows that the outcome would not be very different for a new run with the same setup as our results. This increase in the margin of error is expected, since when all agents learn how to bid and adapt themselves against other co-owners, the success rate can vary more than in the case where the other agents always bid in the same manner.

4.3 Learning Process for Privacy Personas

With the evaluations above, we have measured the success of PANOLA agents in various settings. In these evaluations, the results were given after the agents learn to bid for an amount of content, which corresponds to a training process. With this evaluation, we will investigate how personas learn to perform winning bids from scratch, which can differ because each has differing knowledge and motivation levels. For this setup, we measure the success of a PANOLA agent with a given persona against two *Technician* agents who employ *SBS* that is described in Section 4.1. We have picked a single type persona, *Technicians*, for comparison to reduce the randomness for the opposing agents to showcase the learning curves of each persona better. The reason for picking *Technicians* was because they have a high level of motivation and thus are more active in the auctions that would be more challenging than other personas except fundamentalists. The learning trends are similar against different personas in the same setup; however, the percentage of success when the learning converges differs depending on the persona. Since we aim to show the learning trends of personas rather than the success rates in this evaluation, we omitted the results with all personas except *Technicians* as the opposition. For each persona, the privacy decisions with the PANO auctions will be made for 1,000 pieces of content, and over 50 runs we will demonstrate how quickly they converge to find out the possible winning bids.

According to Figure 6, we can see that the trends match with the outcome of Figure 5, which shows the results after the learning phase. We can also conclude that both knowledge and

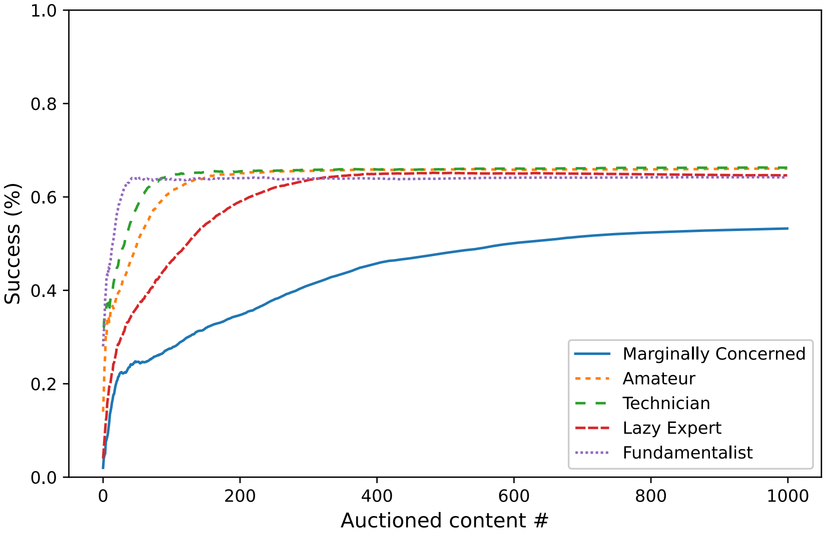

Fig. 6. Learning progress of PANOLA agents with each privacy persona over 1,000 pieces of content.

motivation levels affect the speed of reaching a stable point. Since *Fundamentalists* have the highest knowledge and motivation levels, PANOLA agents that represent this persona are the quickest to converge, followed by *Technicians* and *Amateurs*. When the motivation level is low, as in the case of *Lazy Experts* and the *Marginally Concerned*, convergence requires more auctions since input by these personas is sparse and can be mostly wrong for the *Marginally Concerned* due to their low knowledge level. We can also see that having a higher knowledge level increases the learning speed in the same motivation level when we compare *Technicians* with *Amateurs* and *Lazy Experts* with the *Marginally Concerned*. In summary, all agents quickly reach a certain extent of success, while some personas continue to improve afterward. After about 300 auctions, all personas converge to some level with slight changes over time, with the exception of the *Marginally Concerned*, which even continue to improve after the learning phase with 1,000 items of content. This is to be expected, as users with this privacy persona do not give sufficient correct input to PANOLA agents to build confidence quickly, as opposed to *Lazy Experts,* who also give limited input but mostly correctly due to their higher knowledge level. Therefore, we can conclude that not giving input to the agent when the user is not certain of its own privacy preference would be a better option, since this would let PANOLA agents become more confident about how to perform in auctions.

Table 11 depicts the results of statistical significance analysis with confidence intervals for the learning evaluations. For 1,000 items of content, shown in Figure 6, the statistical significance results show that the confidence intervals are very small (less than 0.5%), with the exception of the *Marginally Concerned*, which has a confidence interval of 5.41%. This is also an expected outcome since our results show that the *Marginally Concerned* still improve after the learning phase, because the feedback received from this persona is low and can be wrong in many cases due to their low knowledge level.

4.4 Employing Values for Individualism or Conformism

In Section 3.1, we explained two parameters in efficiency calculation for bid ranges, namely *Value of Content* (V_{Ct}) and *Value of Conformism* (V_{Cf}). As explained in Section 3.3, a higher valuation of V_{Ct} means that PANOLA agents do not mind the amount of bids they place when content is

Table 11. Statistical Significance Analysis of Experiments Depicted in Figure 6, after 50 Runs per Experiment

Agent Type	Mean of Success	Confidence Interval (95%)
Marginally Concerned	53.23%	±5.41
Amateur	66.05%	±0.31
Technician	66.27%	±0.27
Lazy Expert	64.61%	±0.34
Fundamentalist	64.19%	±0.22

valuable for them, since a higher V_{Ct} enables the agents to consider that winning an auction is more important than spending budget. On the contrary, when the V_{Ct} is lower, agents try to minimize their bids as much as possible in the learning process and therefore the agents who set V_{Ct} low might lose some auctions while trying to find the winning bid. With V_{Cf} valuations, agents set the importance of conforming with the groups in their learning process. With a higher V_{Cf} value, agents conform more with others, leading to a minimization of their taxes.

In this experiment, we demonstrate how these values affect the learning process by measuring the individual success in various setups along with the satisfaction of the entire society by the privacy decisions made. Intuitively, with a lower V_{Ct} and a higher V_{Cf} value, agents should adopt a more conformist behavior, where they might lose some auctions for the sake of society. With the opposite valuations, agents should try to win as many auctions as possible, regardless of the others' privacy preferences. We experiment with two different V_{Ct} and V_{Cf} valuations. For the first one, V_{Ct} and V_{Cf} values are assigned for maximum conformism, meaning that V_{Ct} is set for 0 and V_{Cf} is set for 0.5, while for the second one the agent aims to satisfy its individual goals, with V_{Ct} set for 0.5 and V_{Cf} set for 0. With these two valuations, we investigate two different numbers of co-owners to see the effect of the level of opposition. For the first co-owners setup, we assign three co-owners who have privacy preference opposing the PANOLA agent and only one with the same privacy preference. For the second setup there are two non-learning agents against the PANOLA agent instead of three, making it a two-against-two agents setup. With both co-owner numbers, we investigate the two V_{Ct} and V_{Cf} valuations separately. We have 10,000 items of co-owned content by these agents, and for each they enter a PANO auction sequentially. Similar to our previous experiments, we execute 50 runs for each setup. We use a single *Fundamentalist* PANOLA agent with other non-learning *Fundamentalist* agents, who bid with *SBS* from Section 4.1; however, our results with other persona types are similar.

The **satisfaction of society** (*SS*) is measured by considering how satisfied the co-owners are with the outcome (Equation (7)), where n is the total number of co-owners for the auction and d is the number of co-owners who did not want the outcome. For an individualistic agent, we consider it satisfied only when the outcome of the auction is in its favor. For a conformist agent, additionally, we consider it satisfied when the outcome is different from its actual preference but in favor of at least half of the participants. Since in all experiments the co-owners have conflicting privacy preferences, we expect the satisfaction of society to be closer to the middle percentages, because one's satisfaction with the privacy outcome would not necessarily satisfy another who would prefer another outcome.

$$SS = \frac{n - d}{n} \tag{7}$$

Figure 7 shows the success of individuals as well as the society satisfaction under these four settings. Comparing (a) with (b) and (c) with (d), we can say that the V_{Ct} and V_{Cf} valuations work as expected. With conformist settings, the satisfaction of society becomes higher than the PANOLA agent's own success, while it manages to have higher success than the satisfaction of

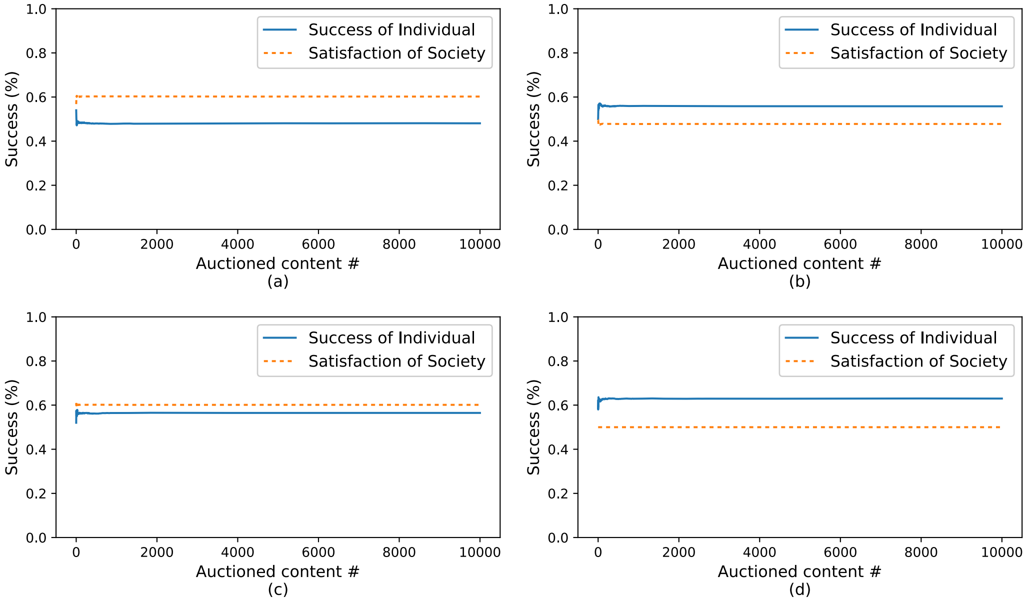

Fig. 7. Success of PANOLA agent and satisfaction of society in (a) one agent on the side of a conformist PANOLA agent against three agents; (b) one agent on the side of an individualistic PANOLA agent against three agents; (c) one agent on the side of a conformist PANOLA agent against two agents; and (d) one agent on the side of an individualistic PANOLA agent against two agents.

society with an individualistic setup. We also note that, since the other agents are non-learning agents, even when the individualistic PANOLA agent has only one other agent supporting its privacy preferences, it can still succeed against three non-learning agents because the PANOLA agent learns how to outbid the others. This experiment addresses RQ-3, where we ask if PANOLA agents can help others preserve their privacy by finding the right balance between individualism and conformism. As a result of this experiment, we observe that when the agents set their valuations high only for the content that is important for them and low for the remaining, they can help others preserve privacy, leading to a higher success in society.

5 DISCUSSION

We have explained how PANOLA agents can assist OSN users in order to protect their privacy for collaborative decisions. We have also demonstrated that PANOLA agents can learn to bid better over time in the PANO auctions. Our experiments show that individuals who employ PANOLA agents obtain a higher success in preserving privacy than those who employ a non-learning agent (Figure 4). The improvement that PANOLA permits is more visible for *Lazy Experts*, the *Marginally Concerned*, and *Amateurs* as their success increases from the $0.1 - 0.3$ range to over 0.7. Thus, PANOLA actually enables their users to preserve their privacy. We observe over many simulations in Figure 5 that, when all users employ PANOLA, the success rate of the users—almost independent of the persona—converges to a stable value, where no persona type is left at a disadvantage. This shows that PANOLA is useful for all persona types. Finally, we observe that if agents can adjust their valuations for some content as to conform with the society, they can help others preserve privacy, leading to a higher success in society. Our simulations assume that the users start with an agent that does not have any previous information about the user at all. In real life, to speed up the learning process, it could be possible to ask the user a few questions up front to elicit his or her privacy perception and then start with a pretrained model that would adapt to the user over time, as is generally done in practical applications of machine learning. In a real-life application, the agent could come with default values for V_{Ct} and V_{Cf} as used in the simulations but could ask the user to update the values according to his or her preferences. Similarly, the agent could also get input from the user about his or her satisfaction with the previous interactions to further improve the learning process. In this article, we have not considered the contextual properties of shared content. In a real-life application, it might be useful to differentiate between a number of contexts as done in the literature [33] so that when necessary, the agents can learn to bid differently in each context. In terms of usage, the agents can act on behalf of the user or can suggest their bids to the user, who could then decide whether to approve it or not. Note that the system does not require every user to use the system in the same way: some of the users might fully delegate the decision to their agents, while other users might choose to bid on their own. Since each agent is responsible for its own user, the different ways of usage by the other agents do not influence the learning process or the flow of the system.

5.1 Related Work

Privacy in ubiquitous systems started to receive attention around the early 2000s, with the Internet becoming accessible to most people in the world and enabling easy sharing of and access to private information over the web. Langheinrich [26] is one of the first studies to investigate the open issues for privacy-respecting approaches for ubiquitous computing. Spiekermann and Cranor [29] and Gürses et al. [18] study the grounds of engineering privacy, explaining how information-related domains can be designed to employ privacy-preserving methods. Paci et al. [28] provide an extensive survey of the literature on access control over community-centric collaborative systems, laying down the key issues and giving a roadmap for future challenges. Bahri et al. [5] show the challenges of preserving privacy over decentralized OSNs and provide a review of previous work done for overcoming these challenges. Bertino and Ferrari [7] discuss the approaches and concepts for applying privacy to big data, which became an essential part of domains such as IoT and OSNs. They lay out challenges to achieve successful privacy approaches for big data domains. These studies all show that privacy is an important aspect of collaborative information systems and address the need for effective mechanisms.

Even though the main goal is intended to satisfy the general good for collaborative privacy decisions, the agents that represent entities naturally have the goal to force their privacy requirements onto others. Therefore, while the environment should be fair to each agent, the agents should have the freedom to try different strategies to be placed in an advantageous position. PANO offers [41] a fair mechanism to decide on which action to take, which uses Clarke Tax auctions at its core with some economic modifications such as group-wise spending, bidding boundaries, and income-expenditure balance levels. For the competitiveness of the agents, we introduce a learning mechanism that is based on reinforcement learning, where agents can adapt according to the visible information resulting from the outcome of previous auctions [40]. We also use an evaluation distance coefficient to overcome the cold start problem for those agents that have no prior information about auctions or their opponents.

Such and Criado [35] focus on the challenges of multiparty privacy in social media, categorize the current approaches to preserve privacy, and present a roadmap for the requirements the multiparty privacy solutions should fulfill. Auctioning is presented as one of the multiparty privacy resolution approaches, and its drawback is given as the difficulties that users can face to understand and manage the process. We tackle this issue in this article with an adaptive agent-based approach that learns to bid on behalf of the user. Thus, the user is never asked for bid values explicitly.

Collaborative privacy management is investigated in the literature on different domains. Fong [17] introduces the **Relationship-Based Access Control (ReBAC)** mechanism and provides a model to make it applicable to OSNs, where users can define their privacy constraints related to the relations that are available in OSNs, such as friends or colleagues. Even though a relationship-based model is suitable for privacy policies that solely depend on predefined relationships, real-life cases are usually much more complicated. The relationship types in commonly used OSNs are usually very limited, and users tend to have policies that include/exclude specific users for all content.

The Multi-party Access Control Model by Hu et al. [19] is another work that focuses on determining a single final policy according to the privacy requirements of the users. It also takes users' sensitivity levels into account and proposes a voting mechanism for the publisher and stakeholders of content. The success of the model is evaluated according to oversharing, undersharing, and correctness metrics. Correctness shows the percentage of correct assignments according to the decisive policy, while oversharing shows the unintended share percentages for content and undersharing depicts the percentage of users where content sharing is intended but not actually performed. We also used these metrics in our evaluation of PANO, and we showed that it indeed performs better than a native Clarke Tax auction approach according to the defined metrics [39, 41]. However, in multi-party access control models, there is no learning of privacy requirements or better formation of a final policy, as we have proposed here.

As an extension to the relationship-based access control mechanism, Klemperer et al. [22] propose using photo tags for defining privacy policies. The main goal of this work is to reduce the complexity of relationship-based policies and take advantage of contextual properties of photo tags. This approach is mainly targeted for sharing of photographs where it is possible to tag the content. Their proposed approach is not meant to be used for collaborative systems, where co-owners need to decide on the final policy. Further, they do not provide a learning component for the tags.

There exist some approaches that use negotiation or argumentation techniques to resolve privacy conflicts between people or software agents in multiagent domains. Such and Rovatsos [36] employ a negotiation-based approach for predefined privacy policy sets of users, and the goal of the conflict resolution is to find a middle ground by negotiation between the agents according to their privacy policies. The approach requires a definition of privacy policies by human interaction, and also the negotiation is still managed by the users of the system themselves. This shortcoming was tackled in Such and Criado [34], and the same approach was extended with modeling and learning user behavior, as well as implementing a software mediator that manages the negotiation process without the need of a human interruption. Another multiagent negotiation model was introduced in Kekulluoglu et al. [20], which includes a comprehensive negotiation protocol to be used by the agents. In addition, incentives of the agents are considered in the approach.

Kökciyan et al. [23] propose an argumentation-based approach for collaborative privacy management in OSNs. In this approach, software agents are employed for representing OSN users according to their privacy requirements and for resolving privacy conflicts where related agents have different opinions for a privacy action. The agents can access the domain knowledge and infer semantic rules that are not directly available as information. Using the argumentation mechanism, agents can attack others' beliefs and assumptions with their own inferred knowledge and aim to convince the other agents to make them accept their own users' privacy preferences. The presented work is promising for those domains where agents can retrieve domain knowledge and infer new semantic rules with limited computational complexity. However, gathering knowledge, inference of information with limited computational power (i.e., memory size, processing power), and ensuring communication between agents for a given time period are some major challenges for several domains such as IoT or widely used OSNs. Therefore, the applicability of the proposed model can become infeasible when the the mentioned limitations affect functionality.

Recently, approaches considering human values and norms for collaborative privacy management have been gaining traction. Calikli et al. [8] 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. Ajmeri et al. [2] study norm emergence factoring in the context of the agents, taking the sanctions into account. In another work, Ajmeri et al. [3] provide a framework where agents aggregate the value preferences of the users and choose ethically appropriate actions for social contexts. Ulusoy and Yolum [42] propose a norm-based access-control mechanism for collaborative privacy decisions, which considers both personal and social norms in decision making. Mosca et al. [27] propose an agent architecture for OSNs, where the agents have essential properties such as explainability and adaptability while being both utility and value driven. Colnago et al. [10] study the IoT domain for personalized privacy assistants, which lays out characteristics of users with a case study in terms of privacy understanding and preferences.

The use of machine learning for privacy is gaining momentum, and the research area is still open for further improvement. Fogues et al. [16] provide an agent-based approach that requires user input when required to learn incrementally about user policies and recommend privacy policies for sharing content for multiuser scenarios. The work differs from ours in the way their system learns the user preferences by user feedback, while in our mechanism agents can learn from the visible properties in the system and the outcome of the collaborative privacy decisions. Vanetti et al. [43] propose a machine learning approach for filtering unwanted textual content in OSNs. The system classifies the texts and learns to prevent them from being published on OSN pages, according to the predefined user requirements. Even though the work is solely based on short texts, the idea can be extended to include different contextual elements for a more generic solution. Squicciarini et al. [31] infer privacy policies of OSN users for photographic content. The policies are generated according to a contextual classification of the images, which are trained with some datasets and user experiences. Zhong et al. [45] employ contextual image properties in a different way: they extract and learn from the image features in a way to detect possible privacy conflicts to take further action. This approach can be beneficial to focus on privately significant content and to exclude non-controversial content from collaborative privacy decisions. Another work by Squicciarini et al. [30] aims to learn privacy features of image content, with a novel approach of employing sentiment in the context of image classification. Fang and LeFevre [15] propose a software wizard for learning the privacy requirements of the users. The approach takes user input and makes use of it to classify users of the system as groups, suggesting similar privacy settings to the users within the same group. Albertini et al. [4] presents another privacy policy learning approach, which is based on ReBAC [17]. The model creates association rules according to the usage history of the OSN users with the Apriori Algorithm and generates privacy policies accordingly. Our approach differs from the previous work in terms of policy generation. In our model, we only require relaxed context-related privacy policies to create agents, and the reinforcement learning process is employed in runtime to fine-tune the policy-related preferences with changing bids. With this approach, the initial input requirement is reduced, and the agents learn the system on behalf of their owners.

5.2 Future Work

In light of this work, some interesting research directions open up. Modeling the opponent is one of these we would like to investigate. We would like to develop an agent that can change its behavior as needed, as well as build models of other agents in the auctions to make better decisions. This can also be followed up by another research question that would investigate integrating social norms for privacy in this process, which we already have been investigating as a separate line of work [42]. This could be beneficial to create learning agents according to their normative behavior that can also befit societal values. Another direction would be to capture the interrelation dynamics between agents, especially those of trust. When agents trust each other, they can reflect this when bidding. For example, an agent might not bid high to share some content when it knows that the other agent would rather not share it, even if it has the budget. This could lead to behaviors where individuals do not act in a self-interested manner but work together to preserve each other's privacy. Another line of work may investigate real-life applicability of this research with user experiments. A user study can be conducted to investigate the levels of interaction and knowledge between the user and the personal assistant, or whether a user can update his or her privacy preferences according to the actions taken by the personal assistant. This can be achieved by developing a user interface for the participants in a real-life social network, to showcase the changes of both users' and agents' privacy understanding over time. These improvements would enable us to employ software agents that assist users in the most efficient way to preserve their privacy in collaborative systems.

ACKNOWLEDGMENTS

We thank the anonymous reviewers whose comments have improved the paper considerably.

REFERENCES

[1] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514.

[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'18).* 22–34.

[3] Nirav Ajmeri, Hui Guo, Pradeep K. Murukannaiah, and Munindar P. Singh. 2020. Elessar: Ethics in norm-aware agents. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems.* 16–24.

[4] Davide A. Albertini, Barbara Carminati, and Elena Ferrari. 2016. Privacy settings recommender for online social network. In *2016 IEEE 2nd International Conference on Collaboration and Internet Computing (CIC'16).* 514–521.

[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] Andrew G. Barto and Sridhar Mahadevan. 2003. Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems* 13, 4 (Oct. 2003), 341–379.

[7] Elisa Bertino and Elena Ferrari. 2018. *Big Data Security and Privacy.* Springer International Publishing, Cham, 425–439.

[8] 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.

[9] Edward Clarke. 1971. Multipart pricing of public goods. *Public Choice* 11, 1 (1971), 17–33.

[10] Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2020. Informing the design of a personalized privacy assistant for the internet of things. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20).* Association for Computing Machinery, New York, NY, 1–13.

[11] Carlos Diuk, Andre Cohen, and Michael L. Littman. 2008. An object-oriented representation for efficient reinforcement learning. In *Proceedings of the 25th International Conference on Machine Learning (ICML'08).* ACM, 240–247.

[12] Jean-Baptist Du Prel, Gerhard Hommel, Bernd Röhrig, and Maria Blettner. 2009. Confidence interval or p-value?: Part 4 of a series on evaluation of scientific publications. *Deutsches Ärzteblatt International* 106, 19 (2009), 335.

[13] Janna Lynn Dupree, Richard Devries, Daniel M. Berry, and Edward Lank. 2016. Privacy personas: Clustering users via attitudes and behaviors toward security practices. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems.* 5228–5239.

[14] Eithan Ephrati and Jeffrey S. Rosenschein. 1991. The Clarke tax as a consensus mechanism among automated agents. In *Proceedings of the 9th National Conference on Artificial Intelligence - Volume 1 (AAAI'91).* AAAI Press, 173–178.

[15] Lujun Fang and Kristen LeFevre. 2010. Privacy wizards for social networking sites. In *Proceedings of the 19th International Conference on World Wide Web (WWW'10)*. ACM, 351–360.

[16] R. L. Fogues, P. K. Murukannaiah, J. M. Such, and M. P. Singh. 2017. SoSharP: Recommending sharing policies in multiuser privacy scenarios. *IEEE Internet Computing* 21, 6 (Nov. 2017), 28–36.

[17] Philip W. L. Fong. 2011. Relationship-based access control: Protection model and policy language. In *Proceedings of the 1st ACM Conference on Data and Application Security and Privacy (CODASPY'11)*. ACM, 191–202.

[18] Seda Gurses, Carmela Troncoso, and Claudia Diaz. 2011. Engineering privacy by design. In *Computers, Privacy & Data Protection*. 25 pages.

[19] 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.

[20] Dilara Kekulluoglu, Nadin Kokciyan, and Pınar Yolum. 2018. Preserving privacy as social responsibility in online social networks. *ACM Transactions on Internet Technology* 18, 4, Article 42 (April 2018), 22 pages.

[21] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Ashesh Rambachan. 2018. Algorithmic fairness. In *AEA Papers and Proceedings*, Vol. 108. 22–27.

[22] Peter Klemperer, Yuan Liang, Michelle Mazurek, Manya Sleeper, Blase Ur, Lujo Bauer, Lorrie Faith Cranor, Nitin Gupta, and Michael Reiter. 2012. Tag, you can see it!: Using tags for access control in photo sharing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'12)*. ACM, 377–386.

[23] Nadin Kökciyan, Nefise Yaglikci, and Pınar Yolum. 2017. An argumentation approach for resolving privacy disputes in online social networks. *ACM Transactions on Internet Technology* 17, 3, Article 27 (June 2017), 22 pages.

[24] Ponnurangam Kumaraguru and Lorrie Faith Cranor. 2005. Privacy indexes: A survey of Westin's studies. 2005. *Available as ISRI Technical Report CMU-ISRI-05-138* (2005).

[25] Airi Lampinen, Vilma Lehtinen, Asko Lehmuskallio, and Sakari Tamminen. 2011. We're in it together: Interpersonal management of disclosure in social network services. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'11)*. ACM, 3217–3226.

[26] 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.

[27] Francesca Mosca. 2020. Value-aligned and explainable agents for collective decision making: Privacy application. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. 2199–2200.

[28] Federica Paci, Anna Squicciarini, and Nicola Zannone. 2018. Survey on access control for community-centered collaborative systems. *Computing Surveys* 51, 1, Article 6 (Jan. 2018), 38 pages.

[29] Sarah Spiekermann and Lorrie Faith Cranor. 2009. Engineering privacy. *IEEE Transactions on Software Engineering* 35, 1 (Jan. 2009), 67–82.

[30] Anna C. Squicciarini, Cornelia Caragea, and Rahul Balakavi. 2017. Toward automated online photo privacy. *ACM Transactions on the Web* 11, 1, Article 2 (April 2017), 29 pages.

[31] Anna C. Squicciarini, Dan Lin, Smitha Sundareswaran, and Joshua Wede. 2015. Privacy policy inference of user-uploaded images on content sharing sites. *IEEE Transactions on Knowledge and Data Engineering* 27, 1 (Jan. 2015), 193–206.

[32] 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.

[33] Anna C. Squicciarini, Smitha Sundareswaran, Dan Lin, and Josh Wede. 2011. A3p: Adaptive policy prediction for shared images over popular content sharing sites. In *Proceedings of the 22nd ACM conference on Hypertext and Hypermedia*. 261–270.

[34] Jose M. Such and Natalia Criado. 2016. Resolving multi-party privacy conflicts in social media. *IEEE Transactions on Knowledge and Data Engineering* 28, 7 (July 2016), 1851–1863.

[35] Jose M. Such and Natalia Criado. 2018. Multiparty privacy in social media. *Communications of the ACM* 61, 8 (July 2018), 74–81.

[36] 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.

[37] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction*. MIT Press.

[38] Ming Tan. 1993. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the 10th International Conference on Machine Learning*. Morgan Kaufmann, 330–337.

[39] Onuralp Ulusoy and Pınar 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.

[40] Onuralp Ulusoy and Pınar Yolum. 2020. Agents for preserving privacy: Learning and decision making collaboratively. In *Multi-Agent Systems and Agreement Technologies*, Nick Bassiliades, Georgios Chalkiadakis, and Dave de Jonge (Eds.). Springer International Publishing, 116–131.

[41] Onuralp Ulusoy and Pınar Yolum. 2020. Collaborative privacy management with auctioning mechanisms. In *Advances in Automated Negotiations*. Springer Singapore, Singapore, 45–62.

[42] Onuralp Ulusoy and Pınar Yolum. 2020. Norm-based access control. In *Proceedings of the 25th ACM Symposium on Access Control Models and Technologies (SACMAT'20)*. Association for Computing Machinery, New York, NY, 35–46.

[43] M. Vanetti, E. Binaghi, E. Ferrari, B. Carminati, and M. Carullo. 2013. A system to filter unwanted messages from OSN user walls. *IEEE Transactions on Knowledge and Data Engineering* 25, 2 (Feb. 2013), 285–297.

[44] Piotr A. Woźniak, Edward J. Gorzelańczyk, and Janusz A. Murakowski. 1995. Two components of long-term memory. *Acta Neurobiologiae Experimentalis* 55, 4 (1995), 301—305.

[45] Haoti Zhong, Anna C. Squicciarini, and David Miller. 2018. Toward automated multiparty privacy conflict detection. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM'18)*. ACM, New York, NY, 1811–1814.

Received July 2020; revised May 2021; accepted June 2021