

Privacy in Collaborative Systems

Onuralp Ulusoy



Utrecht University



SIKS Dissertation Series No. 2022-28

The research reported in this thesis has been carried out under the auspices of SIKS, the Dutch Research School for Information and Knowledge Systems.

© 2022 Onuralp Ulusoy

Cover by Onuralp Ulusoy

Printed by Ridderprint | www.ridderprint.nl

ISBN: 978-94-6458-710-4

DOI: 10.33540/1510

Privacy in Collaborative Systems

Privacy in Collaboratieve Systemen

(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de
Universiteit Utrecht
op gezag van de
rector magnificus, prof.dr. H.R.B.M. Kummeling,
ingevolge het besluit van het college voor promoties
in het openbaar te verdedigen op
woensdag 7 december 2022 des middags te 12.15 uur

door

ONURALP ULUSOY

geboren op 31 augustus 1988
te LULEBURGAZ, Turkije

Promotor: Prof. dr. P. Yolum

Copromotor: Dr. T. Baarslag

Beoordelingscommissie:

Prof. dr. Frank Dignum

Prof. dr. Catholijn M. Jonker

Prof. dr. Albert Ali Salah

Prof. dr. Anna C. Squicciarini

Prof. dr. Jose M. Such

Supervisor: Prof. dr. P. Yolum

Co-supervisor: Dr. T. Baarslag

Assessment Committee:

Prof. dr. Frank Dignum, Umeå University

Prof. dr. Catholijn M. Jonker, Delft University of Technology

Prof. dr. Albert Ali Salah, Utrecht University

Prof. dr. Anna C. Squicciarini, Pennsylvania State University

Prof. dr. Jose M. Such, King's College London



Contents

1	Introduction	1
1.1	An Overview of Privacy	1
1.2	Privacy in Collaborative Systems	2
1.3	Research Directions	8
1.4	Related Publications	11
2	Collaborative Privacy with Auctioning Mechanisms	13
2.1	Introduction	14
2.2	Employing Auctioning Mechanisms Privacy	16
2.3	Agent-Based Bidding	18
2.4	Multi-agent Simulation	25
2.5	Evaluation	28
2.6	Discussion	32
3	Learning Agents for Supporting Users in Preserving Privacy	35
3.1	Introduction	36
3.2	Technical Background	38
3.3	Preserving Privacy with PANOLA	41
3.4	Evaluation	53
3.5	Discussion	64
4	Privacy Values in Collaborative Systems	71
4.1	Introduction	72
4.2	Learning action choices	74
4.3	Effect of openness to change and conservation	76
4.4	Effect of self-transcendence and self-enhancement	82
4.5	Discussion	87
5	Privacy Norms in Collaborative Systems	89

5.1	Introduction	90
5.2	Privacy Norms	91
5.3	PRINOR	93
5.4	Generating Norms	96
5.5	Evaluation	102
5.6	Case Study for Real-Life OSNs	109
5.7	Discussion	112
6	Conclusions and Future Directions	115
	Bibliography	119
	Summary	129
	Samenvatting	131
	Acknowledgements	133
	Curriculum Vitae	135



Introduction

1.1 An Overview of Privacy

Privacy is the notion that characterizes an individual's or a group's ability to keep information about themselves without sharing with others when needed. Individuals or groups can allow selective information to be shared with others, while keeping the rest *private*. While *preserving privacy* seems like a trivial task where everyone can govern the information about themselves, in practice, it is virtually impossible to reach total privacy for everyone since private information can also disseminated by others, either implicitly or explicitly. For example, an individual might want to keep her location private, while a friend of hers can knowingly share it with the public, which would create an *explicit privacy violation*, or can share some other piece of information where the location of the individual might be inferred, which would be an *implicit privacy violation*. Therefore, preserving privacy usually cannot be governed by the individuals alone or groups by themselves, but require elaborate mechanisms that should also consider the actions of others.

Even though privacy was identified as a part of human life since the ancient Greece [29], the systematic definition of privacy was not provided until the late 19th century. Warren and Brandeis published the article "The Right to Privacy" in 1890 [115], which is generally considered as the first publication that advocates privacy [41]. In the article, Warren and Brandeis define privacy as "the right to be let alone", and argues that the term *property* comprises not only tangible forms of possession, but also intangible forms. While by then the privacy of individuals or groups could only be disturbed by the printed media such as newspapers or word of mouth, the technological developments in the first half of the 20th century provided new means to spread information, hence resulting in new ways that can cause privacy violations. Alan Westin was one of the pioneers to shift the discussion from physical forms to embody the privacy into a notion of how the information about individuals are governed. In his book "Privacy and Freedom" 1890 [116], Westin argues that digitization of personal data makes it easily accessible by the government or other entities, and a person should have complete jurisdiction over his or her data. Since the argument is still the basis of privacy now, it can be said that Westin's book laid the foundation for how privacy should be understood. Prosser, who was a legal scholar,

defined four types of harmful activities that can cause privacy violations in 1960 [75]. These activities are listed below.

- Intrusion upon someone's seclusion or solitude
- Public disclosure of embarrassing private facts
- Publicity which places someone in a false light
- Appropriation of someone's name or likeness for gain without her permission

While the time these harmful activities were described is decades before technological tools became a part of our daily lives, it can be said that the same actions are still the main reasons for privacy violations, especially with the internet becoming a massive source of information diffusion. In 2008, Solove published "Understanding Privacy" [85], which overviews the progress made to conceptualize privacy over the years in an interdisciplinary manner, and sets a framework for understanding privacy and offers practical guidance for engaging with relevant issues, such as surveillance, data mining, identity theft, state involvement in information handling, which are still in line with the harmful activities for privacy, listed by Prosser.

Preventing harmful actions that violate privacy of individuals or groups is now a challenge that has been tackled by many disciplines and with many different approaches. However, privacy is "Too complicated to be boiled down to a single essence", as Dan Solove says [85], and resolving privacy issues will stay as an ongoing challenge for the years to come. With the internet being prominent and widely used due to applications such as smart tools, online social networks and collaboration environments, the private information becomes available to multiple entities, and cannot usually be governed by only the ones that introduce the information to a system. Hence, the tools we use result in *co-owned* information, which makes privacy resolution even more challenging. When information is co-owned, it can contain private information of the co-owners, thus a resolution to share a content or not, or whom to share it with, would need to address the privacy concerns of all co-owners. When these concerns differ and co-owners would desire conflicting outcomes for a privacy decision, reaching a resolution would cause some of the co-owners to not be able to have all their privacy concerns addressed, which would require elaborate mechanisms to overcome the privacy violations.

In this dissertation, we will focus on the notion of co-owned information, and will seek ways to develop mechanisms that enable collaboration in a way that preserves privacy of each entity. This chapter will first overview the literature for privacy in information systems, with a focus on social applications where each party can collaborate for preserving privacy. Then, we will lay out our research questions which we will tackle in the following chapters, and will briefly summarize our peer-reviewed publications that were used to shape each chapter.

1.2 Privacy in Collaborative Systems

In the second half of the 20th century, technological improvements enabled organizations to store information about the people as digital data. Naturally, this brought

up the concerns over how the private data should be handled, and the need of regulations to act in line with privacy arose. One of the most influential early efforts to legislate privacy was the "US Privacy Act of 1974" [18], which created the notion of fair information practices to define privacy policies. The principles were based on work of Alan Westin, and consist of seven items, namely *openness and transparency*, *individual participation*, *collection limitation*, *data quality*, *use limitation*, *reasonable security* and *accountability*. Even though the act created awareness over the world where many regulations have been placed over protecting privacy, it was also criticized for slowing processes over businesses. Organization for Economic Co-operation and Development (OECD) codified the fair information practices in their guidelines in 1980 [39], in order to prevent a privacy protection harming the economic growth by creating trade-barriers. Then it took a few decades to take another big step on regulating privacy to protect the rights of the people. In 1995, The European Union's Data Protection Directive [25] was established, which limited data transfers to non-EU countries, where "an adequate level of privacy protection" was required. As for individuals' privacy, the directive also stated an important notion: *explicit consent*. The directive stated that "*Personal data may only be processed if the user has unambiguously given his or her consent*", which still stays as one of the most important pillars of privacy protection.

Until the 1990s, privacy was mostly seen as a regulated information transfer between two parties, and in most cases between the government and their citizens. The end of the 20th century has overseen the rise of the internet, which also changed how privacy was understood. With the internet becoming more and more available for home usage, information exchanges between entities became easier, and the data that contain private information began to be concern of multiple parties. That is, the internet enabled ways to disseminate content for the individuals for not just their own private information, but of the other individuals as well. This change in the tides also placed attention of the scientific researchers to come up with solutions to provide frameworks with which privacy can be regulated. Cranor *et al.* [28] investigates the concerns of the internet users with a survey, which showed that even at the end of the last century, most of the users had concerns about privacy violations while using the internet. The article also emphasizes the importance of standardized privacy policies to regulate privacy, but mentions that many of the users had very limited knowledge about policies and questions the applicability of only policy-based approaches, while suggesting elements of self-regulation and technical approaches to be involved in the process. In 2002, World Wide Web Consortium (W3C) developed *the Platform for Privacy Preferences (P3P)*, which is also a follow up to Cranor's previous studies. P3P aimed to standardize the way for the websites to communicate about their privacy policies and an easy way for the internet users to express their privacy preferences with the help of some tools or software agents [26, 27]. Even though P3P has not reached a wide audience where all the websites adopted its capabilities, it stayed as an inspiration which led to development of other policies in the following years. In the early 2000s, *privacy by design* became an important term, where researchers established principles to follow in designing online systems where data is shared between entities. Langheinrich [57] worked on defining the principles of privacy-aware *ubiquitous computing*. Ubiquitous computing, also named as *pervasive computing* is defined

as computing made to appear anytime and everywhere. Ubiquitous computing can occur using any device, in any location, and in any format, which is applicable to many systems such as wearable devices, mobile phones and computers with the help of the internet. Langheinrich defines a comprehensive set of guidelines for designing privacy-aware ubiquitous systems, which he places into six main items, namely *notice, choice and consent, anonymity and pseudonymity, proximity and locality, adequate security* and *access and recourse*. Cavoukian [21] defines seven foundational principles for *privacy by design* in networked data systems and technologies. These principles are listed as below:

- *Proactive not Reactive; Preventative not Remedial* approach in order to prevent privacy violations before occurring.
- *Privacy by default*, which indicates that the initial choice should always be to protect privacy of the individuals.
- Privacy should be embedded into design.
- Privacy should not be considered as a counter-part to security, and rather be considered a way to improve both.
- Full lifecycle protection, which means that in every stage of information transfer, privacy should be taken into consideration.
- Visibility and transparency over how the data is processed.
- Making design choices that are user-centric to respect user privacy at all times.

Gurses *et al.* [44] tackle *privacy by design* from an engineering perspective, and describe the previous work to have vague definitions that can be interpreted in different ways, which can cause issues in handling privacy. The article describes two case studies to show how privacy principles should be applied in each step of a design of engineering systems, and argues that generalizable methodologies that build upon the principle of data minimization are essential to properly address *privacy by design*. Spiekermann [86] describes *privacy by design* as the notion that holds the promise to solve the privacy issues of digital system, but mentions that there are immense challenges to make it fully applicable. The article lays out the main challenges for *privacy by design* as privacy being a fuzzy concept, which is usually confounded with security and is very difficult to protect. Not having an agreed-upon methodology to design privacy related practices is one of these difficulties. Another difficulty listed by the article is having little knowledge about the tangible and intangible benefits and risks associated with companies' privacy practices when designing their products. Spiekermann also mentions that in order to have a generalized approach a guideline such as European Data Protection Directive [25] should be followed, or a new guideline should be established which takes the current scale of internet technologies into account. As suggested by this article, The General Data Protection Regulation (GDPR) was introduced by European Union (EU) in 2016 and enforced for use in information systems in 2018. Since then, the GDPR became an important component of regulating privacy in EU which addresses the transfer and use of data within and

out of EU. The main goal of the GDPR is to enable individuals to control and have their rights over their personal data, and to simplify the regulatory environment for international business. The GDPR, which supersedes the European Data Protection Directive, was also adopted by many countries outside EU and is currently the most established privacy regulatory system to protect the privacy of individuals or groups. The principles of GDPR states that personal data can only be used if there is a legal basis to do so. These lawful purposes for GDPR are given as below:

- If the data subject has given consent to the processing of his or her personal data;
- To fulfill contractual obligations with a data subject, or for tasks at the request of a data subject who is in the process of entering into a contract;
- To comply with a data controller's legal obligations;
- To protect the vital interests of a data subject or another individual;
- To perform a task in the public interest or in official authority;
- For the legitimate interests of a data controller or a third party, unless these interests are overridden by interests of the data subject or her or his rights according to the Charter of Fundamental Rights (especially in the case of children)

While the lawful purposes can enable the use of personal data, the GDPR also states the conditions clearly to avoid ambiguity. According to the GDPR principles, consent for data use must be explicit for data collected and each purpose data is used for. In addition, consent must be a specific, freely-given, plainly-worded, and unambiguous affirmation given by the data subject; by an online form which has consent options structured as an opt-out selected by default is a violation of the GDPR to avoid accidental consent by the users. Data subjects also must always have the right to withdraw their consent at any time, which should be as easy as opting in to data usage. The GDPR currently offers the most efficient way to protect privacy of individuals while they use online systems governed by companies, where each user has the means to preserve their privacy. However, the systems where information does not pertain to a single individual and instead a group of people, each person in the group should have a say in giving consent, which would not be easily applicable with the GDPR. For example, a user of an online social network can make use of GDPR to give consent over data they share, and the OSN provider would process the data accordingly. However, the same person can share data which depicts others, as in a group photo, and the GDPR does not enforce the OSN provider to obtain consent from the other people that are depicted in such content, which might cause privacy violations by making the content public. Hence, preserving privacy still is a challenge that needs to be tackled for the systems that contain *co-owned* content. Moreover, it should be stated that there is not an agreed definition of what privacy is due to it being a plural [84] and contestable [71] subject, which makes developing privacy resolution mechanisms even more challenging.

In computer systems, when completing a task or achieving a goal is done by more than a single entity, it is defined as *collaboration* [60]. Therefore, the systems that

incorporate collaboration become *collaborative systems*. Collaborative systems encompass the use of computers to support coordination and cooperation of two or more entities who attempt to perform a task or solve a problem together [11]. The capability to collaborate with users and enable collaboration between the users is essential for assisting users in finding the information they need and solving the problems they have [42]. Tolone *et al.* [103] is one of the earlier survey studies in the 2000s that investigates collaborative systems in access-control, focusing on requirements and models. Zhu [121] focuses on role mechanisms for collaborative systems. Paci *et al.* [74] provide an extensive survey on privacy in community centered collaborative systems up to recently, which aims to categorize previous research for collaborative privacy in terms of the ways to handle privacy policies and the evaluation methods. Since privacy for co-owned content involve more than a single individual on reaching a resolution, collaborative systems approaches become essential, where each involved party plays a role on preserving privacy.

When privacy decisions involve more than a single group or individual, handling these decisions is defined as *multi-party privacy*. Multi-party privacy has been a hot topic for researchers over the last three decades, with the internet providing many applications that enable collaboration in social systems, thus having the need of handling *co-owned content*. One of the early but influential work is by Sandhu *et al.* [78] in 1996, which introduces *role-based access control* in handling privacy in a community. In this work, roles and role hierarchy play an important role to decide on who can access which type of information. A role represents the set of permissions needed to carry out a certain job function, while a role hierarchy describe the structure, where some roles can inherit the permissions of other rules due to being higher or lower in the hierarchy. The users do not have individual privacy policies, and instead roles are assigned to the users which define their privacy settings. Thomas and Sandhu [101] extend this work to propose *task-based access control*, which assigns the permissions to access pieces of information with tasks instead of roles. Thomas [100] presents another alternative with *team-based access control*, which is a direct extension to *role-based access control*, but assigns roles to groups of people instead of the individuals.

In the early 2000s, online social networks (OSNs) started to emerge, which allowed individuals to connect with each other with defined relationship types (e.g., *friendship*) over the network, form groups and share content either publicly or with a set of people within the OSN. While OSNs gained popularity, the immense amount of users, connected with various relationships and sharing tremendous amount of content required new ways to handle privacy preservation. Lederer *et al.* [58] stated that managing roles or groups for privacy in an OSN is a significant burden for the OSN users, which results in not being able to manage sharing content in a privacy preserving way. Jones and O'Neill [49] investigate the feasibility of group-based privacy control in social networks and aim to automate the groups with clustering algorithms. However, having full automation for group forming in OSNs stays impossible due to the highly dynamic nature of the system, where a tremendous amount of users form and dissolve groups all the time, and even within the groups contextual properties play a big role on privacy decisions [73].

To overcome the drawbacks of early access-control models that aim to formulate

certain privacy rules, researchers provided many alternative solutions that can be employed for protecting privacy in OSNs. Gates [40] introduces *relationship-based access control*, which aims to manage privacy according to interpersonal relationships within OSNs. Fong [38] represents privacy constraints as formulas in modal logic, which enables OSN users to define their privacy requirements accordingly with the relationships present in the OSN, such as friends or colleagues. Squicciarini *et al.* [88] propose an auction-based model where users enter auctions for deciding on a policy that requires collaborative management over a content. On another line of work, Squicciarini *et al.* [90] develop a privacy manager, which aims to produce semi-automated privacy rules for online social network users and fill the gap between the privacy requirements of the users and the privacy protection mechanisms that OSN providers offer. Gurses and Diaz [43] argue that researchers treat different scientific problems for privacy in online social networks independently, while the problems such as surveillance, institutional privacy, or social privacy for OSNs are entangled and could benefit from a more holistic approach. Such and Rovatsos [95] employ a negotiation based approach where privacy policies of users are pre-defined, and the privacy resolution is reached by negotiation between the users according to the sets of privacy policies they own. Such and Criado [93] extend this work by modeling and learning user behavior within the OSNs, where a software mediator handles the negotiation process without the need of human input. Kökciyan *et al.* [54] make use of argumentation techniques to develop an argumentation based approach for collaborative privacy management, where software agents represent OSN users and use an argumentation mechanism to convince other agents to make them accept the privacy preferences of the users they represent. Misra and Such [64] develop a personal assistant agent that recommends personalized access control decisions for social media users, based on the social context and utilization of users' social media profiles. Ratmajer *et al.* [76] propose a variation of the one-shot Ultimatum Game, and model users interacting with each other to reach privacy decisions about shared content. Ilija *et al.* [48] develop a collaborative multi-party access control model that makes use of OSN users' social relationship where collective privacy policies can be applied with trusted connections. Ajmeri *et al.* [2] incorporate social norms for agent-oriented software engineering methods to develop socially intelligent personal agents that are privacy-aware. Mosca *et al.* [70] provide an agent architecture that considers explainability, role-agnosticism, adaptability while being utility and value-driven to achieve collaborative privacy management in OSNs. Mosca and Such [67] extend this work with an explainable agent that supports multiuser privacy, aiming to identify optimal sharing policies and justifying the optimality with argumentation-based explanations.

Learning how to handle multi-party privacy is another important aspect that researchers focused on recently. Akcora *et al.* [6] adopt an active learning approach for risk estimation from user interactions in social networks, and aim to associate a risk level with OSN users to provide them information about how risky it would be to share content with others. On another line of work, Akcora *et al.* [7] develop an approach to learn user similarities on OSNs by a novel similarity measurement which considers network and user profile similarities. Kökciyan and Yolum [53] provide a semantic approach for learning to detect privacy violations in social networks by performing reasoning over contextual privacy information. Fogues *et al.* [37] develop

a computational model to predict an appropriate sharing policy for a given scenario by building a classifier based on machine learning. Tonge *et al.* [104] propose a method that combines object and scene-based tags for social media content which are uncovered using convolutional neural networks to learn and predict privacy for sharing images over OSNs. Singh *et al.* [83] aim at a privacy-aware design for sharing personal data with third parties by applying a semi-supervised learning method.

1.3 Research Directions

Privacy is a notion that is synchronous to the dissemination of information. When a piece of information is disseminated, it becomes known by other entities. What we deem private can only be known by others if it is shared with others in a form that can depict information, such as digital data, physical sources (e.g., written texts) or communication between people. With the internet technologies dominating our daily lives with new applications such as collaboration via cloud computing or interaction via online social networks; the ways to share information became so accessible that even the least tech savvy people are able to use many of these applications. With the worldwide use of these applications, information that people consider private are uploaded to the internet and governed by either the service providers, or by the users themselves with the provided privacy settings within the applications. This raises several challenges that needs to be tackled for the preservation of privacy. The users of online applications might not be able to express their own privacy requirements, resulting in unintentional violation of their own privacy. An another challenge is that a piece of content can depict information for more than a single individual (e.g., a collaboratively edited document, a group photo etc.), and other users that have access to the same content can disseminate it without the permission of each party that are depicted in the content. To overcome these issues, a comprehensive mechanism that provides the users the means to reflect their privacy is essential. Such mechanism should also ensure collaboration, since the content in online applications usually contain more than a single individual, making them *co-owned*. The collaboration in the mechanism should provide equity to each co-owner so that a selective part of the community cannot lead the privacy decisions. Moreover, the mechanism should prevent the absolute governance by the service providers in a centric manner, and enable distributed computation as much as possible.

Figure 1.1 lays out a scenario depicting a group photo to be shared in an online social network, and the effects of this decision on people's privacy. As seen in the example, sharing a group photo on social media may affect the privacy of several individuals. Widely used social media platforms leave the share decision to the individuals who upload the photo. However, as stated in our example scenario, individual decision can lead to privacy violations, or a shared piece of content to not reach its target audience. This example shows the necessity of a collaborative decision, which should consider all individuals that can have their privacy affected by an outcome. While providing such a mechanism to online social network is essential, resolution of privacy disputes would need to consider several aspects to satisfy privacy requirements of all individuals. For example, a simple mechanism such as voting to share



Figure 1.1: *An example scenario depicting the challenges of collaborative privacy in social media.*

or not share would not be able to capture a comprehensive decision, since it would again result in privacy outcomes that differ significantly from individuals' privacy preferences, which might include items such as sharing with a selective audience, or conditional decisions depending on the opinions of others and the society itself.

In this dissertation, we address these challenges to provide a privacy resolution mechanism that enables collaboration, distribution of tasks, use of software agents to assist users and make use of human values and social norms in privacy decisions. To establish our mechanism, we select online social networks (OSNs) as our domain, and aim to make privacy decisions that preserve privacy of OSN users while sharing content such as group photos or videos over the network. We evaluate our results with multi-agent simulations, case studies and indicate the open research questions that can serve as future directions to improve the proposed mechanisms. Here, we list the main three research questions we tackle in this dissertation, and will explain our work in detail in the following chapters to show how we address these questions.

Research Question 1: *What is an easy-to-use privacy resolution mechanism that enables collaborative privacy decisions in online social networks?* Currently, widely used social networks do not provide the means for collaborative decisions, and content that can affect more than a single user can be shared with the sole decision of the user

that uploads it. In an ideal setting, each user that has private information should have a say in the privacy decision, making her a *co-owner* of the content. The privacy decisions should be reached collaboratively by all co-owners, and the mechanism should offer each user a fair chance at making the outcome in line with her privacy requirements. The most popular OSNs are used by millions of people world wide, and content shared over OSNs reach tremendous amounts. Having a complex mechanism that takes time and effort would be inapplicable for real-life cases, therefore a simplistic approach in collaborative privacy decisions is important. The mechanism should also be easy-to-use, since not every OSN user would have a high level of knowledge to establish their privacy requirements in the decisions. We envision a mechanism where software agents can act on behalf of users and these agents would benefit from performing in an easy-to-use approach that requires less deliberation. Another essential feature is that the privacy resolution mechanism should reach decisions with little computation and with as few iterations as possible.

Research Question 2: *How can we design software agents that assist online social network users on their privacy decisions?* While a collaborative privacy decision mechanism that can be used by OSN users would provide the means to preserve privacy, requiring input for each privacy decision by each co-owner of a piece of content can become a tedious task, and still result in privacy violations when some users do not put the required effort into decisions. Providing assistance to users is essential for privacy decisions, which can be achieved with software agents. If the privacy requirements of the users can be captured with agents, they can learn the expectations of the users and assist them in the collaborative privacy decisions, reducing the need for user input significantly. Moreover, without the use of agents, the users of OSNs can have different levels of knowledge about privacy, or can have differing motivations in their decisions, which would result in some users dominating the privacy decisions and causing privacy violations for others. Hence, privacy assistant agents should provide equity for all types of users and aim to learn the privacy requirements of even the least knowledgeable users with no motivation to provide input.

Research Question 3: *How do we identify human values and social norms in privacy decisions, and incorporate them in the privacy decisions?* While people make decisions in their lives, they are lead by *human values* they adopt, which affect the choices they make. In privacy, these values can still be in effect, and can play a role in the collaborative privacy decisions made. Since human values are not mutually exclusive for individuals and people share some character traits, similar privacy behavior can be expected for the OSN users that share the same values. These shared actions result in social norms to emerge in the social systems, where in a given context a group of people behave the same. OSNs are social systems where we would expect social norms to emerge, and privacy decisions can also be in line with these *social norms*. If these norms are identified and incorporated in the privacy decision mechanisms, they can provide resolutions to privacy disputes just by users conforming to norms. Employing norm-based decisions can also reduce the need of a complex mechanism where each user needs to state her privacy requirements, reach a decision where some conflict resolution might also be needed. Instead, the identified norms can be prompted to the users in a given context and if they are willing to comply,

privacy decisions can be taken accordingly with the norms without the use of another mechanism.

In light of these research questions, we propose our mechanisms that aim to provide solutions to the challenges stated in this dissertation. In Chapter 2, we tackle the first research question, and provide a mechanism called PANO, which offers an auction-based method based on Clarke-Tax auctions, where OSN users can bid for privacy actions in order to decide sharing content over the OSN or keeping them private collaboratively. Chapter 3 focuses on the second research question, and proposes software agents that called PANOLA, which aim to learn the privacy requirements of the users they represent and assist them in collaborative privacy decisions to preserve their privacy. Finally, Chapters 4 and 5 tackle the third research question. In Chapter 4, we investigate the effect of human values in privacy decisions, and infer ways to make use of them in privacy decisions. Chapter 5 builds on top of this to provide norm-based privacy decisions with a mechanism called PRINOR, which identifies emergent social and personal norms for privacy, and employs these norms in privacy decisions.

1.4 Related Publications

In this section, we will list the scientific publications that were published with parts of the work in this dissertation. The following four chapters will describe these works in detail, while making connections between each work to create a complete approach to tackle collaborative privacy resolution.

Chapter 2, which focuses on providing an auction-based collaborative mechanism, contains the work from the publications below:

- Onuralp Ulusoy. 2018. Collaborative privacy management in online social networks. In Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, 1788–1790.[106]
- Onuralp Ulusoy and Pinar Yolum. 2018. PANO: Privacy Auctioning for On-line Social Networks. In Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (Stockholm, Sweden) (AAMAS '18), 2103–2105.[107]
- Onuralp Ulusoy and Pinar Yolum. 2020. Collaborative Privacy Management with Auctioning Mechanisms. In Advances in Automated Negotiations. Springer Singapore, 45–62. (Best Student Paper Award) [112]

The first publication lays out the challenges of collaborative privacy management and the roadmap to follow for collaborative privacy mechanisms. The second publication states the basis of PANO mechanism, and the third publication describes the details of the mechanism, along with the evaluation of its performance in preserving privacy in OSNs.

Chapter 3 extends PANO mechanism with the introduction of software agents named PANOLA, which learn privacy behavior of OSN users and assist them in collaborative privacy decisions. The following publications are contained in the chapter:

- Onuralp Ulusoy and Pinar 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.[110]
- Onuralp Ulusoy and Pinar Yolum. 2021. PANOLA: A Personal Assistant for Supporting Users in Preserving Privacy. *ACM Transactions on Internet Technologies* 22, 1, Article 27 (sep 2021), 32 pages.[113]

The first publication provides a learning mechanism for OSNs where software agents can learn and act on behalf of OSN users. The second publication extends this work by considering privacy personas of the OSN users, and offers a refined mechanism with personal privacy assistants.

Chapter 4 investigates the human values in terms of privacy, and evaluates how OSN users with different values perform for collaborative privacy decisions. A part of the evaluations in this work is published in the first publication for Chapter 3 above.

Chapter 5 builds on top of Chapter 4, and introduces privacy norms that are employed for a mechanism that identifies social norms and makes use of them in privacy decisions. The following publications are used to form this chapter:

- Onuralp Ulusoy and Pinar Yolum. 2019. Emergent Privacy Norms for Collaborative Systems. In *PRIMA 2019: Principles and Practice of Multi-Agent Systems*.3690 Springer International Publishing, Cham, 514–522.[108]
- Onuralp Ulusoy, Pinar Yolum. 2019. Privacy Norms in Online Social Networks. In *Benelux Conference on Artificial Intelligence*[109]
- Onuralp Ulusoy and Pinar Yolum. 2020. Norm-Based Access Control. In *Proceedings of the 25th ACM Symposium on Access Control Models and Technologies (Barcelona, Spain) (SACMAT '20)*. Association for Computing Machinery, New York, NY, USA, 35–46. [111]
- Albert Mwanjesa, Onuralp Ulusoy, and Pinar Yolum. 2022. DIPP: Diffusion of Privacy Preferences in Online Social Networks. In *Advances in Social Simulation*. Springer, 29–40.[72]

The first paper explains a norm-based mechanism where emergent social norms for privacy can be identified and applied for privacy decisions. The second paper is a brief summary of the norm-based privacy approach for OSNs. The third paper extends the first and second work with a wider variety of norms, and improves the mechanism with a hierarchical approach where agents have the ability to conform or ignore social norms depending on the privacy requirements of the users they represent. The last paper investigates privacy behavior diffusion in OSNs, and lists the significant factors that affect a social norm forming or becoming out of date.

Collaborative Privacy with Auctioning Mechanisms

ABSTRACT

Collaborative systems contain pieces of information pertaining to multiple individuals, such as a picture of a group of friends or a collaboratively edited document. When this information is to be shared, it becomes a challenging task of how to share this information and with whom.

Online social networks are a prime example of collaborative systems, which enable users to share content with other users. Many times, a shared content, such as a group picture, may reveal private information about the uploader as well as others who are associated with the content.

Ideally, protection of privacy in such cases would need to consider the privacy concerns of all relevant individuals. However, these concerns might conflict and satisfying one user's privacy needs could cause a privacy violation for others. This calls for computational mechanisms that can decide on the privacy policies of the content collaboratively.

In this chapter, we propose an agent-based collaborative privacy management model for Online Social Networks (OSNs), namely PANO. In PANO, agents represent OSN users and manage their privacy requirements on their behalf. We extend Clarke-Tax mechanism [22] for auctioning to achieve fair handling of privacy settings and to tax the agents whose privacy settings are chosen.

PANO is designed as an auction mechanism where OSN users can be represented by software agents, which can bid on behalf of the users considering their privacy concerns as privacy policies. For a piece of content, all the associated users are represented by agents, and bid for privacy decisions on either to share or not share the content or whom to share it with. We aim to provide a mechanism that prevents abuses so that all users have a say in the decisions. We expect the decisions to prevent privacy violations, while still enabling sharing content over the OSNs when it is possible. We evaluate our approach over multi-agent simulations to measure how PANO copes with collaborative privacy decisions.

2.1 Introduction

Collaborative systems enable users to interact online while sharing content that pertains to more than one user. 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, an IoT system, in which one security camera would like to share footage of a setting to guarantee security for the people, while one individual would prefer to keep the location of herself secret. In both of these cases, the content being 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 the uploader and hence is actually *co-owned* by others [46, 95].

Reaching a consensus in collaborative privacy management is usually not an easy task, since people's privacy requirements can easily be in conflict. In real life cases, the resolution requires much time and effort through mechanisms like negotiation, argumentation, and so on. When co-owners have different privacy requirements, they should be given the means to make a decision as to either share or not to share the content. However, current systems enable only the uploader to set privacy settings while publishing contents, but does not allow co-owners to state their constraints. Since they can't provide collaborative solutions, people tend to resolve conflicts via different media [56] or in most cases, ignore others' privacy requirements and cause voluntary or involuntary privacy violations. In order to be able to handle them online, decision mechanisms should be in place.

Ideally, systems should provide privacy management mechanisms to regulate how content will be shared. Recently, multi-agent agreement techniques have been used to address collaborative privacy management. Kekulluoglu *et al.* [50] and Such and Rovatsos [95] propose negotiation-based approaches that enable users to reach a consensus on how to share a content. Kokciyan *et al.* [54] use argumentation to enable one user to persuade the other into sharing with her own privacy constraints. These approaches have been successful but require heavy computations; that is, they can only be used when the entities can reason on its privacy policies and communicate with others intensively. For example, a negotiation approach would require each party in the decision process to value varying privacy settings for deciding which one would be a better outcome for a party. Since privacy has many implicit or explicit aspects such as the relationship between the parties, the time and location associated with content, contextual properties such as a business meeting, a family party and so on, placing a value on privacy decisions can be a very challenging task for negotiation. Furthermore, any privacy decision can require many iterations while all the parties reevaluate their stance in the negotiation, which would be time consuming and infeasible for the OSN domain, where the amount of content shared is tremendous. Argumentation approaches share similar drawbacks when all contextual properties are considered, and reaching a consensus might be impossible in many cases, since there is no ground truth for some privacy decisions. An example for this could be that a person can argue that each picture taken for a business meeting should be shared among colleagues while another person would never wants to be associated with that business and considers each picture in a business meeting private. In this case, these

two parties would never reach a consensus regardless of how many iterations are made to reach a decision. Considering these drawbacks, another approach where it would be easier to value privacy decisions and with less consumed time, preferably a one-shot approach, can be more fitting for the OSN domain.

On a different line, Squicciarini *et al.* [88] propose a model where users enter auctions for deciding on a policy that requires collaborative management over a content. Each user creates bids based on how much she wants to see a content public or private. In that approach, users collect units by publishing content and tagging people that are related to content. These units are used in an auction, where users spend their units to convince other users to accept their policy, based on Clarke-Tax mechanism [22, 34]. However, this system is open to abuse by the users, such that a single user's privacy can be ignored repeatedly when all others collaborate strongly. Further, since the bids are expected to be generated by users individually for each post, it is difficult to apply them in real-life online social networks.

Ensuring collaborative privacy management in a real-life scale OSN requires the system to scale to the tremendous amount of content being shared. To enable this, first, the operations expected from users should be handled automatically so that users do not need to think through the operations for each content. OSNs enable users to share a great amount of content, and each user can also have some privacy concerns over the content that others share, such as a group picture including the given user. Asking the user her opinion about every single decision would be time consuming and most of the OSN users would not be willing to spend this much effort. Second, the proposed mechanism should be easy to compute, because it will be repeated for each content separately. If each decision taken by the mechanism requires long runtime to process, or requires iterations, it would result in a heavy overload for the OSN providers, since the well-known OSN applications host millions of new shared content each day. Third and most importantly, it should preserve privacy of the users fairly so that no user is left at a disadvantage. In a mechanism that offers resolutions for collaborative privacy conflicts, the users who have more knowledge or motivation than the average users can take advantage of the mechanism to enforce their own privacy preferences. A fair mechanism should prevent that and aim to ensure each user, regardless of their knowledge or motivation levels, has a similar amount of effect in the privacy decisions.

Accordingly, in this chapter we propose PANO, an agent-based collaborative privacy management system that uses ideas from the work proposed by Squicciarini *et al.* [88]. There are three main contributions of PANO: First, it employs agents for privacy management, where agents act on behalf of users to enforce their privacy constraints, so that heavy user involvement is reduced to minimum. The agents manage their users' privacy constraints and bid on behalf of them. Second, it contains a fair reward mechanism, which is protective against abuses, and at the same time encourages users to share content online. Third, it works with a group-wise budget system in auctions, where the agents cannot use the advantages they gain from the system against individuals. This disables an agent to abuse another agent's privacy. Our experimental evaluation shows that agents can indeed carry out this task successfully and help preserve their users' privacy with high accuracy.

2.2 Employing Auctioning Mechanisms Privacy

As a broad definition; privacy is the concept of individuals deciding on how much about themselves to be shared with the others. In OSNs, these decisions can be represented with privacy policies. Applying privacy policies when the information is solely related to an individual itself is an easy task, when the necessary tools are provided. However, a piece of information, e.g. a photo content in an OSN, can be related to more than one individual. In such cases, the decisions of the individuals for the extend of how much to share may differ, resulting in conflicts. These conflicts require some resolution mechanism to define a generalized privacy policy with the goal to comply with every individual’s privacy requirements. For this conflict resolution mechanism, we propose PANO, an agent-based collaborative privacy decision system, where agents employ auctioning mechanisms to reach decisions on privacy conflicts [107, 112]. PANO employs an extended version of Clarke-Tax Mechanism[22] with an agent-based approach as the underlying mechanism.

2.2.1 Background: Clarke-Tax Mechanism

Clarke-Tax mechanism provides an auction mechanism, similar to English auctions where participants bid for different, possible actions in the environment. The participants earn units for their actions and can bid with those units for actions. The action that receives the highest total units from the participants wins and is executed. Different from an English auction, participants who aid in the winning action to be chosen, i.e., that bid towards it, are taxed according to the value they put on it. This is achieved by subtracting the bid values of every single user from the overall values. If the subtraction of a single user’s bid changes the overall decision, it shows that the user’s bid on this action had a *decisive* value. Thus, the user is taxed with the difference of the actual action’s total units and the units of action to be taken if that user was not present in the auction.

The Clarke-Tax auctions are beneficial for decision making for multiple participants with different opinions, as they support truthfulness [88]. If Clarke-Tax auctions are applied in commerce, then each participant would have their own budget (e.g., money) to bid with. However, since we are emulating the auction idea, the participants are given budgets 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 her current budget.

Table 2.1: *Four User Bids for Sharing an Image*

Users	No Share	Limited Share	Public Share
Alice	3	5	0
Bob	15	2	0
Carol	5	8	5
Dave	2	6	18

In the context of collaborative privacy, Clarke-Tax mechanism is used to decide on how an image is going to be shared. Squicciarini *et al.* [88] consider three types of

sharing actions: *no share*, *limited share*, and *public share*. We follow the same scheme here. When an image is about to be shared, all the relevant participants bid on these three possible actions. Table 2.1 shows an example of biddings of four users for deciding to share or not share a content. Users decide based on their own importance of the three actions. According to Table 2.1, it can be seen that Bob values the *No Share* action more than the others, while Dave values *Public Share* action the most.

According to the biddings of all users, Clarke-Tax auction mechanism decides on which action to take. Based on the bids from Table 2.1, *no share* action receives a total of 25, whereas *limited share* receives a total of 21, and *public share* a total of 23 points. Therefore, the *no share* action is chosen. Table 2.2 shows the resulting decision and applies the taxes according to the biddings in Table 2.1. According to the decided action, each user that bids for the decisive action is taxed. In this case, Alice and Bob are taxed as depicted in Table 2.2, since each of these users' absence in the auction causes the decisive action to be changed. When the units of Alice is subtracted from the overall unit totals, the decision of sharing the content receives the maximum of 23, while not sharing gets a unit total of 22. This causes Alice to be taxed with a unit amount of 1. As mentioned above, Bob bid a greater value for not sharing (15 units), and its absence from the auction also causes the final decision to be changed. Since the differences of the bids for different actions are much bigger in Bob's case (i.e. 13 units, obtained from the subtraction of *Public Share* and *No Share* unit totals), Bob is taxed with a greater value. It is important to note that although the user is taxed, he gets the action to be decided what it values most, and prevents the content from being shared.

Table 2.2: *Clarke-Tax Mechanism Example - Decision and Taxes*

Values	No	Limited	Public	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

The importance of 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. This is due to the bids which change the final decision being taxed by the mechanism, by the amount of effect they make on the outcome. If one bids more than its truthful evaluation, it would cost more and result in having less budget for the future auctions. On the contrary, if Bob had bid for much less, then his privacy might have been violated since the bid might not be enough for his choice to be favored. A lower bid than the correct and truthful evaluation can be costly to lose some important auctions for the participant. Hence, it is important to be able to create bids that reflect the true evaluations of the users.

2.2.2 Challenges

Auctioning with Clarke-Tax Mechanism is an efficient way of negotiation, since it has been shown that truthfulness is the best strategy for bidding [88]. The bidder who overvalues a decision to get her way can be taxed with a greater amount, because it changes the group decision by spending system budget more than the other participants of the auction. This results in the participants bidding with truthful values, while trying to establish its own decision and not get taxed with a greater amount. Even with the notion of truthfulness, applying a pure Clarke-Tax approach still has some limitations that can result in abuse by the bidders or inflation in the budget used. We aim to prevent this with the use of some limitations in the budgets. Consider the following examples:

Example 1 Dave wants to share a photo in which he is with Alice, Bob and Carol. Alice and Carol do not want the picture to be shared. Dave has previously shared a lot of photos that are unrelated to Alice, Bob and Carol and gained substantially more budget than the others. Since Dave can spend more than the others, he puts a high bid (e.g. > 50 for the bids in Table 2.1), instead of the 18 that was bid for sharing in Table 2.1, to share the photo, which bids of Alice, Bob and Carol cannot match. Therefore, the photo is shared in the social network due to Dave being able to use the budget he earned with the others previously, even though Alice, Bob and Carol did not want the photo to be shared.

Another drawback of the application of Clarke-Tax mechanism over OSNs is that it requires user involvement for every single auction. This could be necessary for some cases, but it could become a tedious task for greater number of contents. Also, unavailability of some users in auctions who are related to a content that has privacy conflicts could make the method get stuck, or decide on a semi-successful policy. To resolve these issues, an automated auction process where software agents represent the users and act on their behalf can be implemented.

Example 2 Alice, Bob, Carol and Dave want to decide on sharing or not sharing a photo they are in, in a social network. Carol shares a lot of photos daily, has a lot of budget to spend, and does not want to join a Clarke-Tax auction, but still wants the photo to be shared. Alice also doesn't mind if the photo is shared or not, and doesn't enter the auction. Bob and Dave have a disagreement, and join into an auction, and Bob wins the auction for not sharing the photo, by outbidding Dave. The final policy is decided for not sharing the photo, due to Alice and Carol not participating in the auction. If all were participating, Carol would have placed a bid for sharing the photo, which would have changed the outcome in favor of Carol and Dave instead of Bob.

2.3 Agent-Based Bidding

The pure Clark-Tax based mechanism in Squicciarini *et al.* [88] does not show how participation can be automatized if user involvement for auctions is necessary, which could become a tedious work for the OSN users that share a multitude of contents every day. Thus, we develop an agent-based approach, where each user is represented

with an agent that maintains its user's privacy constraints, manages total budgets, and generates bids when necessary. In principle, understanding users' privacy constraints automatically is difficult. It would require the user behavior to be modeled and privacy constraints to be learned over time. There is a good body of literature on learning privacy constraints [91, 93]. In this chapter, we assume that the user's agent is already aware of the constraints, either through learning or through elicitation, while we will delve into learning user requirements to define the privacy constraints in the next chapter.

2.3.1 Privacy Policy

PANO depends on the decisions of the participants. These participants have a self-contained evaluation calculation to decide on the importance of different actions, and bids according to the result of the evaluation. If an agent can assess the importance of a content for its user correctly, and bid with neither excessive nor low amounts, it can help to get better decisive actions, while preserving the previously obtained budget of the user within the network.

PANO makes use of policies for the agents to compute the bidding evaluations. Agents have multiple policies that correspond to different actions, and in an auction, they correspond to these policies to place bids accordingly.

Definition 1 *PANO Policy:* A PANO policy P is a 5-tuple $P = \{a, n, p, q, i\}$, where a is the agent that the policy belongs to, n is the audience of the policy who are the users affected by the outcome, p is the contextual properties for the content that the policy will be applied, q is the privacy related action and i is the importance of the policy, which is a rational value between 0 and 1.

An example policy of an agent that represents Alice, who wants to share a group photo that was taken in a family picnic with her family members and friends can be defined as:

Example 3 :

$P = \{Alice, \{family[Alice], friend[Alice]\}, photo[Picnic, Family, Outdoor], share, 0.9\}$.

As seen in Example 3, a is defined as the agent of Alice, while the audience of the policy n is family and friends of Alice in the domain, which are affected by the action q of this policy, that is to share the photo. The contextual properties of the policy (i.e., p), are picnic, family and outdoor, which makes this policy valid for every photo that are in this context. Finally, the importance value i is given as 0.9, which means Alice considers this policy an important one and wants all photos in this context to be shared with the policy's audience. Since Alice can have many policies, from which some can have some common properties with this policy, the importance value becomes useful to prioritize between policies and avoid conflicts. For example, if another policy has some shared properties with this policy but have a *no share* action and a lower importance value, the agent of Alice can place the policy in our example more important and try to achieve it more than the conflicting one.

2.3.2 Preventing Abuses in the Auction Mechanism

The Clarke-Tax auctions are beneficial for decision making for multiple participants with different opinions, as they support truthfulness. However, the economic system and the budget used in the mechanism can allow abuses, as explained in Section 2.2.2. In order to prevent the system from facing malicious behaviour by some users, some modifications are need to be made for earning the units and spending it. For privacy, the malicious behavior for auctions could be that some users cooperating together to dominate the collaborative decisions on their behalf, aiming to have a bigger budget than others in order to impose their privacy preferences in all auctions or making use of the lack of knowledge and motivation of other participants to lead the auctions. The main modifications proposed in this chapter to prevent abuses are the group-wise spending, boundaries of the bids for the auctions and the balance of the budget.

Group-wise Spending: To prevent abuse of using units for irrelevant auctions with different users, earned units can only be used in new contents with same co-owners. Recall that there are three types of actions: not sharing with anyone, sharing with a limited number of people, and sharing with public. Limited number of people is decided from the conflicted share decisions of users. For example, when a co-owner of the content decides not to share it with a specific user and another co-owner wants to share the content with it, this user is added to the limited audience list. Since the conflicts, and the audience lists for actions are only related to the policies of the co-owners, spending pre-owned budget from previous auctions with different co-owners would give some participants an unfair advantage. This would result in some users cooperating to share trivial co-owned contents between themselves, and not spend any units for the auctions since the contents have no shared value at all. This is prevented with group-wise spending, where the budget earned from auctions with some co-owners can only be spend in the future auctions with the same co-owners.

Boundaries: Clarke-Tax mechanism allows users to bid as much as the budget they hold. This free market approach economy adds a level of uncertainty to the auctions, since a participant cannot have a clear opinion about what others might bid. A person that has big earnings can bid with high numbers, and make little of the taxes, since they can spend more than the others. Limitations to minimum and maximum bids allowed can be beneficial to prevent users that are richer in the budget from dominating the decisions. This also helps agents that participate in the auctions to have better evaluation functions, because they can have a better opinion about the other participants' bids, especially with prior knowledge about the others' characteristics and the context of the content. Therefore, we enforce a maximum bid boundary in PANO.

Balance: With the notion of minimum-maximum boundaries, the balance between budget earnings and expenditures come into consideration. Users earn units by being a co-owner of a content, and spend them for decisions in auctions and taxes. The amount of units given to each user for each auction is a fixed value, which should be decided according to the characteristics of a domain. For collaborative privacy decisions, earning units equal to or more than the maximum expenditure for each auction would cause the economy to bloat and inflation to emerge. Therefore, earnings from being a co-owner should be less than one can spend for an auction of a content.

We propose a balance where the units earned from a content should be half of the maximum boundary of an auction. Another benefit of this approach is to encourage agents to spend more wisely, since spending less for a relatively unimportant decision could help agents to spend more in the future decisions.

Considering the group-wise spending and boundaries, we define a PANO auction as below:

Definition 2 *PANO*: PANO auction is defined as a 6-tuple:

$AUC = \{c, AC, A, m, M, BD\}$, which consists of the auction's related content c , a set of privacy actions (AC), the set of agents (A) that participate in the auction, minimum possible bid (m), maximum possible bid (M) and the set of placed bids (BD), where each bid $b_{t,a}$ ($b_{t,a} \in BD$) is related to one single action t ($t \in AC$) and one single agent a ($a \in A$).

Given a PANO auction defined as in Definition 2, a system can compute the outcome for the agents, and update their budgets accordingly. At the end of each auction, each participant is given an amount that is equal to the half of the maximum possible bid. This prohibits the agent to bid for the maximum possible bid for each auction. That is, the agent would need to save its acquired budget for the next auction to be able to bid higher than average possible bid. Our reason to employ this half of the maximum boundary is that if an agent acquires more budget than she should use, she would be able to bid the maximum allowed amount for every auction. In this case, it would not make sense for an agent to deliberate the bid amount, since a higher bid would increase her chances to force the action she wants, regardless of the significance of the action. On the extreme opposite case, if the agents would earn very little amount for every auction, they would not be able to bid for many decisions when they consider the content sensitive. In this situation, many privacy violations might occur, and agents would be forced to save their budget for many cases to be able to have a decision in one. One of the most common privacy violation examples is defined as *oversharing*. *Oversharing* in privacy can be described as when a privacy decision results in a content being accessed by the parties that are not part of the intended target audience [122]. While preventing *oversharing* is important, the contrary cases can also create the issue of *undersharing*, which can be defined as some of the intended target audience not being able to access the shared content. Even though *undersharing* cannot be considered as a privacy violation, it is not desirable by both OSN providers and their users, since one of the main goals of OSNs is to help reach various types of content to their target audiences. Our decision to give half the amount of the maximum possible bid aims to find a balance between these two extreme cases, where agents should deliberate about placing their bids to be able to enforce their decisions only when necessary, but they would still be able to enforce their decisions in the sensitive cases, if they bid reasonably. With this approach, we aim to find a middle ground while privacy violations such as *oversharing* are prevented, while the shared content still reaches their target audience, thus averts *undersharing*.

Consider Example 1 where Dave could spend 50 on a picture because he could afford it. With the proposed scheme, group-wise spending would only enable Dave to spend pre-owned budget earned from co-owned contents with the same group.

Therefore, the notion of total owned budget becomes obsolete, replaced with group-wise spending. In this situation, all the co-owners would earn the same amount from the co-owned contents, which would provide equality in auctions. Dave can only bid as much as the other co-owners, so he cannot take advantage by using earned budget with contents co-owned with other people. The agent based policy system can also solve the minority dominating the majority problem in Example 2, which is caused by the absence of co-owners in the auction. We call this a *non-participation problem*. In this example, Alice and Carol, who do not join the auction which results in the only person, Dave, that doesn't want the photo to be shared impose his policy on the others. With PANO, every user in the OSN is represented with an agent, which facilitates privacy policies to bid on behalf of their users. Therefore, even if the users are not available for an auction, their agents always are, bidding according to their owners' policies. With the help of automation, Alice and Carol will also be represented in the auction with agents. Since Clarke-Tax mechanism is based on truthfulness, with every co-owner present in an auction, the resulting action will be decided according to everyone's opinion, instead of a minority of the co-owners.

Example 4 Dave wants to share a photo he is in with Alice, Bob and Carol. Alice and Carol don't want the picture to be shared, and they want to do a Clarke-Tax auction to decide on the final policy. Dave shared a lot of photos before, that are unrelated to Alice, Bob and Carol, and gained bid units from those contents. Since the mechanism only allows Dave to use units obtained from the contents that are related to all Alice, Bob, Carol and Dave, the units from unrelated content can't be used to gain advantage. The outcome of the auction is decided over who values their decision most, with the use of only the gained units from the previous shared contents of the same group, without any unfair advantage to any participant.

2.3.3 Bidding Mechanism

Agents bid on auctions based on their privacy policies. As explained in Section 2.3.1, policies have importance values. On top of this, agents also have privacy characteristics, which is the notion of how much privacy-aware an agent is, represented with a value between 0 and 1, named as characteristic coefficient. For a content in an auction, an agent checks its related policies and determines the set of social network users that it wants to share or not share the content with. The characteristic of the agents and the importance of related policies determine how much the agent wants to bid for an auction, according to the given actions of these policies. In PANO, each type of action can receive a different bid from the participants, so even when an agent has conflicting policies (i.e. different policies of an agent preferring different actions for the same network users), it can place bid on conflicting actions, according to their evaluations. Agents should also consider how much budget they own, and place their bids accordingly (e.g. bidding small amounts when short on budget or bidding higher when have enough spendable units).

In PANO, the bidding mechanism is a linear function that returns an integer value between the bidding boundaries for each action, namely *share*, *no share* and *limited share*, explained in Section 2.2.2. The function considers privacy policies, agent characteristics, the current unit balance of the agent and the number of users

that are in conflict against other agents in the auction to compute the outcome. The function is executed as below:

- For a content, an agent first considers its related policies to find *coefficient* of biddings. This coefficient (PC) is calculated as the mean average of the importance values of the policies.
- To spend the obtained units in a prudent manner, we consider both sharing and not sharing actions at the same time for the mean average, if the audience of both policies are the same. That is, we consider not sharing importance as negative, and sharing importance as positive values. For example, if an agent has two policies for a content; one sharing with importance of 0.8 and one not sharing with the importance of 0.2 for not sharing it with the same audience, the importance coefficient is calculated as the mean average of 0.8 and -0.2 , thus 0.3 for sharing the content.
- This coefficient is multiplied by the privacy characteristic coefficient (CC) of the agent, in order to take into account how privacy aware the user that the agent represents is.
- The final parameter is calculated according to the bidding boundaries and the current unit balance of the agent, namely spendable budget (SB).
- If the agent has two times of the maximum boundary (mb) as the balance, the bidding is made with the multiplication of the maximum boundary and the computed coefficient. If the budget is less, the coefficient is multiplied with the agent's owned budget (ob) divided into two, in order to save budget for the future auctions.

The formula for generating bids for sharing or not sharing is:

$$SB = \begin{cases} mb, & \text{if } ob/2 \geq mb \\ ob/2, & \text{otherwise} \end{cases}$$

$$BidValue = \sum_{\alpha=1}^n (P_{\alpha}\{i\})/n \times CC \times SB \quad (2.1)$$

In Equation 2.1, n is the number of related policies for an auction and $P_{\alpha}\{i\}$ is the importance value for policy α . Intuitively, an agent generates a high bid if the privacy importance of the image in question is high, the agent values privacy (CC) and it has the resources to spend on the bid (SB).

In addition to biddings for sharing or not sharing a content, an agent can also bid for the *limited share* action. *Limited share* action aims to share the content with the conflicting set of users, *i.e.* the users that at least one agent is willing to share the content and at least one agent does not want to share the same content with. For bidding to share with a limited audience, the agent checks how many of the users in the conflicting set of users it wants to share the content with, and simply calculates the ratio of the users fitting into its policies divided by the total number of users in the

conflicting users set, namely as fitting ratio (FR). The agent bids half of the maximum boundary or half of its own units if it has less, multiplied with the computed ratio for the conflicting users and its characteristic. The formula for computing the bid values for limited sharing is given in Equation 2.2, where n is the count of conflicting users considered for the policy, and β is the count of users in the conflicting list fitting into user's own share policies.

$$BidValue(LimitedShare) = \beta/n \times CC \times SB \quad (2.2)$$

2.3.4 Metrics

The success of the mechanism depends on how the final policy out of an auction satisfies the policies of the agents. The resulting policy of an auction should correctly assign the policy-applicable users of the network where there are no conflicts between the auction participants, and try to assign the rest as satisfactory as possible to protect the common good. Equation 2.3 measures how well the overall result found with PANO satisfies the n agents that enter the auction. Success is defined as the number of the users that the applied policy differing from the agents requirements (UPC : count of the users with *unsatisfied policy* for user u), divided by the entire set of users that were considered to share the content with (TNU : total count of users in auction participants' network). An *unsatisfied policy* defines the cases where the action of the policy is not applied to the audience of the policy. Referring back to Example 3, if the outcome of an auction is to share a photo within the context of picnic, family and outdoor only with the family of Alice, it would mean that the policy is unsatisfied for the friends of Alice, since the policy also indicates the photo to be shared with Alice's friends. It should be noted that Equation 2.3 considers each satisfied or unsatisfied policy to contribute to the success in a linear manner. While in reality, the effect of a privacy decision on different recipients could be dissimilar, where a share decision can affect a person's privacy significantly while having a less drastic effect on others. We will investigate the effect of contextual significance in the following chapters, while we left this notion out in this section in order to focus on the performance of PANO.

$$Success\% = \left(1 - \frac{\sum_{u=1}^n UPC}{TNU}\right) * 100 \quad (2.3)$$

Example 5 Consider two agents in a network of 200 agents that enter an auction to decide how to share a picture. The first agent wants to share the picture with 140 people and the second agents wants to share it with all. Assume that as a result, the picture is shared with 160 users that include the 140 users that the first agent preferred. The metric would result in $(1 - (20 + 40)/200) * 100 = 70\%$ success.

With the given policy notation, satisfaction of individual users can also be calculated. Equation 2.4 measures the user satisfaction after an auction, considering how well the outcome is aligned with the agent's policy and the importance of the policy. That is, while the satisfaction value for a single content can be computed with the Success metric, making use of importance values of policies can give us sensitivity

levels (SL) of users for conditions of content types that are also represented in the policies. For example, if a user has importance level of 0.6 for a policy that is related to a condition of a content type, it can be assumed that the sensitivity level of the user for the same condition can also be given as 0.6. For multiple policies of the same user for the same condition, we take the average of importance levels of the related policies for SL value. Using the Satisfaction metric for a single content as CS, and the sensitivity level of the content for the user as SL, we define the User Satisfaction (US) metric for an agent with the formula below, where i is the content id from the previously policy applied contents.

$$US = \frac{\sum_{i=1}^n (SL_i * CS_i)}{\sum_{i=1}^n (SL_i)} \quad (2.4)$$

Example 6 Alice wants to share a photo tagged with "bar", which also has other co-owners. She has a policy $P = \{Alice, \{, friend[Alice]\}, photo[bar], share, 0.7\}$, which indicates photos that are tagged with "bar" to be shared with her "friends", with a sensitivity level of 0.7. According to the final policy decided with PANO, the content is shared with 70 of her friends, while Alice had 100 friends that were involved for the given policy. The satisfaction of Alice for this content is 0.7. If Alice have shared some content before with a satisfaction of 0.6 where her policy had sensitivity level of 1.0, the combined user satisfaction is calculated with the User Satisfaction metric as:

$$US_{Alice} = (0.6 * 1.0 + 0.7 * 0.7) / (1.0 + 0.7) = 0.64$$

2.4 Multi-agent Simulation

In order to evaluate the success of PANO, a multi-agent simulation environment is implemented, where agents represent users in a social network with privacy policies, and use PANO auctions to decide on the share policies of co-owned contents. The policies are defined according to rules that regular social network users tend to rely on, and the simulation checks the success of final policy decisions with PANO auctions over a chunk of co-owned contents.

2.4.1 Context-Based Privacy Constraints

Social networks users have privacy constraints for the contents they share in the network. These constraints can be represented by rules, which then can be used to model behavior of the users. Kekulluoglu *et al.* have done a user study to extract privacy rules of frequent users of online social networks [50]. In this work, there are seven decisive rules, which most of them had contextual content properties as a constraint. Two out of these seven rules were eliminated from the current work, since the simulation doesn't support location or event aspects of the social networks. The remaining rules, with relation types and contextual constraints are presented below:

1. If the user :x is included in a photo that depicts a depressed mood, don't share the photo.

2. If the user :x is included in a photo that is located in a bar then don't share the photo with family members
3. If the user :x is included in a photo that is located on a beach don't share with work related people.
4. If the user :x is included in a photo that is a mature content, don't share the photo.
5. Do not share photos with user:y.

In our policy notation, we define contextual constraints with conditions combined with content types. In the example policy given in Section 2.3, *photo[scenery]* depicts that the content type is a photo, and it's categorized as scenery. We represent the photo categories presented in the rules above with such keywords, which can be obtained with automated photo tags. In PANO, we assume that the photo contents are already tagged, so we can assign them into policy conditions. The conditions can have multiple content types, and also multiple constraints for these content types, giving us flexibility in defining the policies. Since the rules in Kekulluoglu *et al.* [50] focus on not sharing contents in specific conditions, we would require some sharing rules that can create conflicts we aim to resolve in scope of this work. Therefore, we generate some opposing rules to the ones above, where some agents would like to share the photos in the same conditions. For example, when an agent has the third rule in their policy, we can generate a policy for another agent about sharing photos tagged with beach with its friends. In this example, if both agents are co-owners for a photo on a beach, the intersection of OSN users from work related people for *Agent #1* and friends for *Agent #2* will be in a conflicting situation, because the former do not want to share the photo while the latter wants to share it.

2.4.2 Execution of the Simulation

PANO is implemented in Java. The simulation represents social network environments, with the aspects of users, contents, different relation types between users and agents. Agents have defined rules, based on the combinations of relations and contextual tags of photos, as explained in Section 2.4.1.

As mentioned in Section 2.3, agents are only allowed to use the units they obtained from earlier shared contents from the same co-owners. Another prevention for abuses is that each agent can only spend a total maximum of 20 units for each content, so that none of the agents can overvalue it to force its own action. According to the characteristics of the agents, they can disperse their units to different actions, and they can even decide not to spend any units if they don't value the sharing or not sharing of the content. The exact values for bidding boundaries and earned units in our simulations are shown in Table 2.3. Referring back to the balance mechanism we discussed in 2.3.2, each agent is given 10 units, which is half of the maximum bid boundary (20), in order to make the agents consider their decisions carefully and spend wisely for the outcomes that are in line with their privacy policies. We also give each agent an initial 100 units for their first auction, which aims to prevent a cold-start problem where the bid mechanism might not be understood to its full

Table 2.3: *Boundaries and Units Given to Agents in PANO simulations*

Minimum-Maximum Bid Boundaries	0 – 20
Initial Units Given for First-time Users	100
Units Given for Each Auction	10

extent by the users. For each auction, all the bids are collected from each user in a manner where each one does not know the bids or decisions of the others, and the collected information is only known in the system. Then, the simulation calculates the overall decision and taxes of the agents (if any), and updates their budgets within the network, as explained in Section 2.2.1.

Algorithm 1: Main process of the simulation

```

1 load the social network;
2 set relationships;
3 add contents with contextual tags;
4 assign co-owners to the contents;
5 add user characteristics and rules;
6 foreach Content do
7   foreach Co-Owner do
8     determine the set of people to share or not share;
9     if the decision for a person is differing then the other co-owners then
10      | remove person from other lists put the person on the conflicting list
11     else
12      | if the person has no share decision & the person not exists in any
13      |   lists then
14      |   | put the person on no share list
15      |   if the person has share decision & the person not exists in any lists
16      |   then
17      |   | put the person on share list
18   foreach Co-Owner do
19     | place bids according to the final lists
20 trigger PANO;
21 update budgets of the co-owners;

```

Algorithm 1 explains the main process of the simulation. Lines 1-5 initialize the environment by loading the network, setting relationships between the users, adding contents with co-owner assignments and generating agent characteristics and rules. Then the simulation enters into a loop for each content that requires PANO auction for deciding the final policy. This loop defines the share, *limited share* and *no share* lists of a content in lines 8-15. Using the formulas of bid calculations in Section 2.3, agents place bids on each action in lines 16 and 17, and the last two lines update units according to the outcome of the auction, considering the placed bids and the taxes.

2.5 Evaluation

The main goal of a privacy management system is to provide the means to each user to express their privacy requirements and come up with resolutions that fit these requirements as much as possible, even with conflicting policies of several participants. We evaluate PANO using the overall system success and user satisfaction metrics presented in Section 2.2.2.

2.5.1 Correctness of Policies

The first evaluation of PANO is aimed to find out the success of the proposed modifications of the method to the native Clarke-Tax mechanism, to prevent abuses and create more satisfactory privacy policies to the community. The success metric defined in Section 2.3 is used to measure the correctness of policies.

Hypothesis 1 *Given a set of agents with various characteristics, PANO can make privacy decisions more successfully than the native Clarke-Tax approach.*

Using the success metric, PANO is compared to the pure Clarke-Tax mechanism approach, where group-wise spending or boundaries are not present. The evaluation is made with five randomly generated social networks of 1000 users, and 200 randomly generated photo content with contextual tags related to the rules presented in Section 2.4.1. We define *starting points*, which depict the time PANO is introduced to the OSN. An earlier *starting point* means the mechanism was introduced while the OSN is formed and not much content was shared, while a later starting point means the users in the OSN already shared a significant amount of content before we start to measure. Each network was calculated with considering different starting points where some of the contents are considered shared before. These starting points are decided as 40, 80, 120 and 160 contents shared before, respectively. For representing the pure Clarke-Tax mechanism, two *levels of abuse* were used in comparisons to model the examples in Section 2.2.2. First level, which we call *limited abuse*, enables some agents to bid more than the defined 0-20 range, but resulting in consuming all budget if used for every auction. Also, the agents that abuse to protect contents from sharing or try to disseminate to their entire social network are balanced, so the abuser agents can also have conflicts. Second level of abuse is defined by giving unlimited budget to the abuser agents, and only letting one-sided abuse for a single auction, (i.e. having only agents that want to share the content or doesn't want the content to be shared). With Hypothesis 1, we aim to measure if PANO can perform well regardless of the starting point and the level of abuse by the agents, while we expect the higher abuse level with unlimited bid units to decrease the performance of the native Clarke-Tax mechanism. After running the simulations with these networks for three different models, the average of satisfactory policy metric for five networks, according to different starting positions is shown in Figure 2.1. It shows that with PANO, correct assignments of users with created policies increase, with approximately 95% success rate. Limited abuse level with pure Clarke-Tax had approximately 88% success rate, and unlimited budget for approximately 76% success rate.

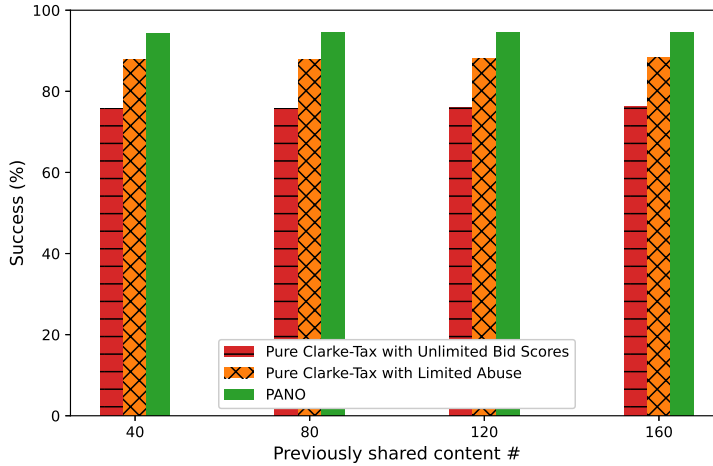


Figure 2.1: Percentages of the social network users that are correctly assigned within the final policies according to co-owners' characteristics.

The same setup is also used for comparing oversharing and undersharing percentages of created policies between the two models. Oversharing covers the users a content is shared with, when some co-owners do not want to share the content with those users. In opposite, undersharing covers the set of users that some co-owners want to share the content with, but not shared with this set of users. The results of oversharing and undersharing evaluations are shown in Figure 2.2 (a) and 2.2 (b), respectively. For oversharing metric, PANO and pure Clarke-Tax method with limited abuse almost performed the same (3.5%), with pure Clarke-Tax approach performing slightly better (less than a half percent). They both performed better than pure Clarke-Tax method with unlimited budget for abuser agents, which had about 5.6% oversharing percentage. The similar performances in oversharing is also investigated, and it has been seen that the abusing agents have more advantage with not sharing the content, since deny-overrides is the default strategy for equal situations. Deny-overrides denotes that when two decisions have the same amount of units for the bid, the decision to not share takes precedence over the decision to share. In some cases, limited abuse level had two conflicting abuser agents trying to outbid each other with the agent that prefers not to share always winning when the same bids are made, and oversharing percentages managed to be as low as PANO. For the undersharing metric, this causes both abuse levels with pure Clarke-Tax method to perform far worse than PANO, where PANO had about 1.5% undershare while two levels of abuse methods had about 8.8% and 18.5%, respectively.

Overall, PANO performs better than applying pure Clarke-Tax mechanism with both abuse levels, considering all the different starting points. Since the evaluated contents were same for all starting points, the success rates were quite similar in four different setups. Even with the same contents, there is a slight increase in the performance when previously shared content number increases, indicating that with large number of pre-shared contents, PANO will perform better.

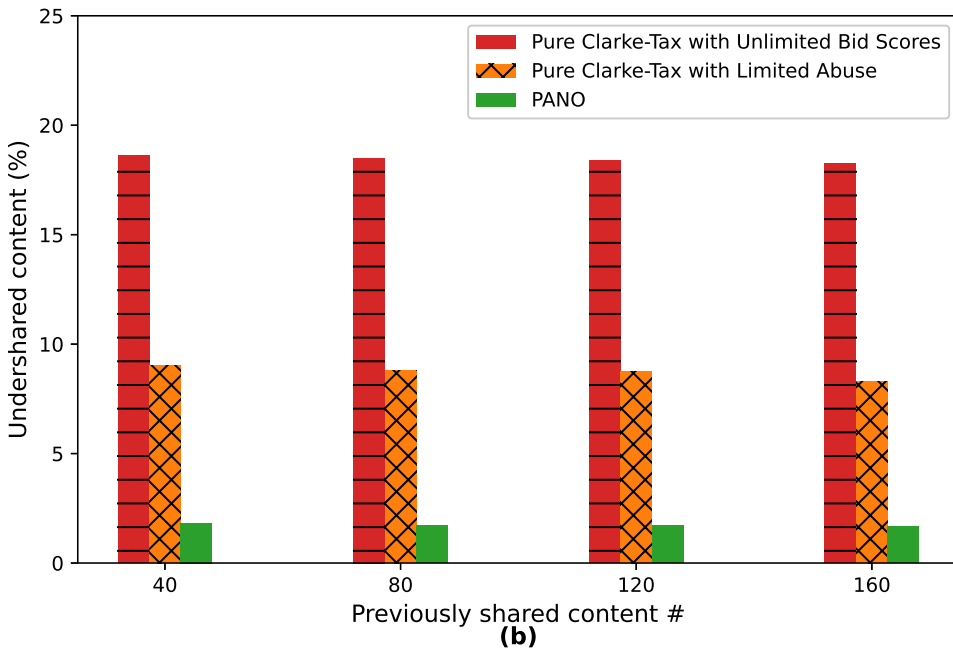
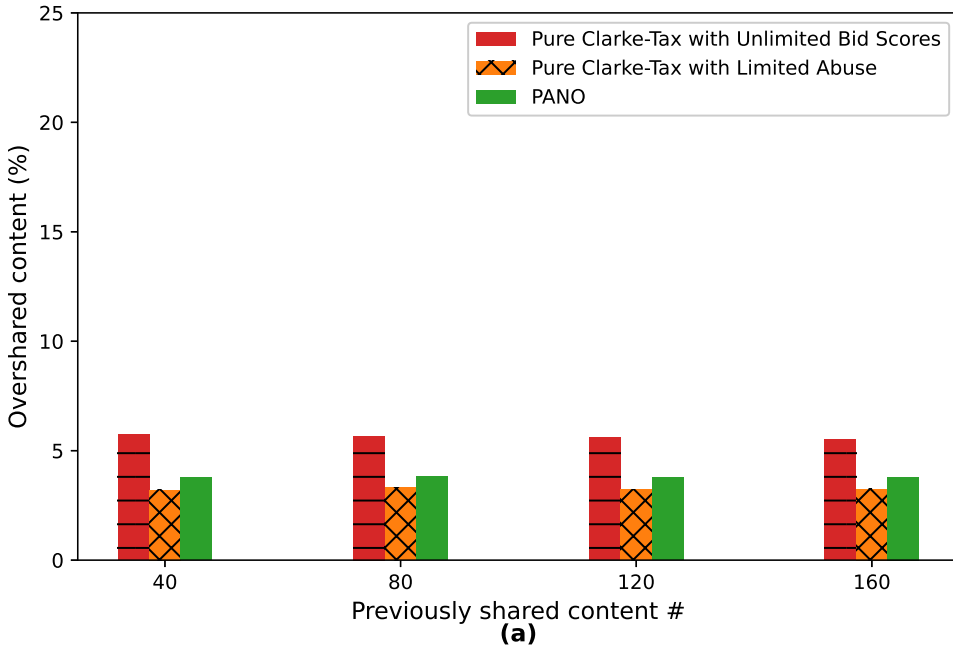


Figure 2.2: a) Oversharing percentages of the social network users b) Undersharing percentages of the social network users

2.5.2 User Satisfaction in Case of Conflicts

The second evaluation setup aims to calculate the user satisfaction metric described in Section 2.3 in a larger generated network. The main goal is to create conflicts between users for a large number of contents, and calculate ratios of intended/unintended audience for decided policies of these contents, using PANO and comparing it with the results of the application of a pure Clarke-Tax mechanism.

Hypothesis 2 *Given that there is a conflict of policies between agents for every auction, PANO will ensure the user satisfaction of most of the agents.*

Since the privacy requirements of all users do not always align, conflicts for the privacy decisions can occur. In a mechanism where some agents can abuse the system to make decisions in favor of only a part of the community, the overall satisfaction of the users within the system would decrease. Since one of our main goals is to prevent abuses with PANO, we hypothesize that we can ensure the majority of the community with it.

The user satisfaction metric in this evaluation has three measurement types. First is the satisfaction ratio, which shows the percentage of people that the policies of the co-owners matching the final policy, taking the sensitivity of the initial policies of the agents into account. The second and third metrics are the percentages of oversharing and undersharing, which are in the same concept as the first evaluation, but these also include the sensitivity levels of the agent policies.

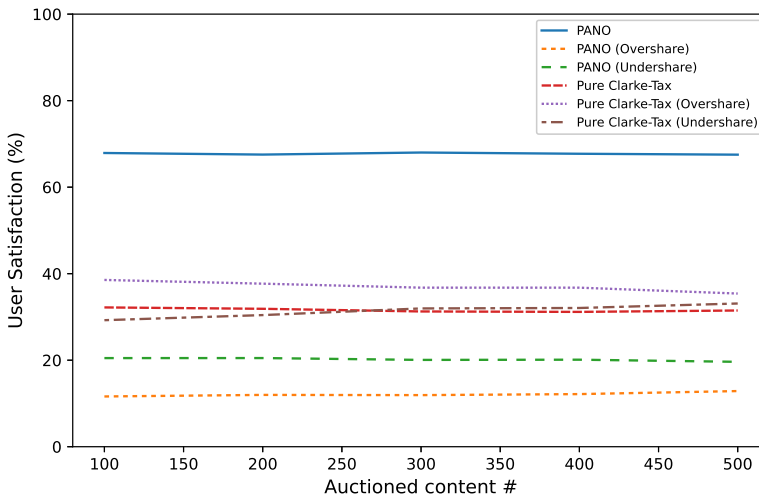


Figure 2.3: Comparison of user satisfaction for various configurations reveals PANO outperforms others in user satisfaction.

The setup of the second evaluation is constructed with a 1000 agent network, where 100 agents owned contents and had approximately 250 friends each. 500 contents were generated with tags, and randomly assigned to 4 of the 100 agents within the network. Policies of the agents were also randomly generated with different sensitivity

levels. For each content, at least one conflicting policy were assured with additional policy generation for some of the co-owners, to create 1 against 2, 1 against 3 or 2 against 2 conflicts. The experiments were repeated 10 times for statistical significance. We compare PANO with the pure Clarke-Tax approach, where at least one agent is trying to abuse the mechanism to get decisions for sharing or not sharing the content according to its own policy. The user satisfaction, oversharing and undersharing metrics were calculated at each 100 content mark according to equation (1), and these three comparisons are presented in Figure 2.3. The results show that PANO outperforms the pure Clarke-Tax approach significantly for user satisfaction due to preventing abuse and enabling the privacy decisions be in line with most of the users' privacy policies. While achieving this, PANO also reduces both undershare and overshare percentages, which means PANO prevents privacy violations better than the pure Clarke-Tax approach in the presence of privacy conflicts and still allowing more content to be shared in the OSN when there is no privacy concerns involved.

2.6 Discussion

Applying multiple privacy preferences from several users for a single OSN content to define a single privacy policy is a major challenge, and it can even be more complicated when these preferences are conflicting. Several studies worked on collaborative privacy management, applying different methods with different approaches. Role Based Access Control, which was applied widely to other software or operating systems is not fully suitable to OSNs, since it can't capture binary or context-dependent relationship features that are available in most OSNs. Fong [38] introduces 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. ReBAC is one of the earliest works which included user related privacy constraints in access control decisions in OSNs, and provided a basis for many collaborative privacy management mechanisms.

Such and Rovatsos [95] propose a model, where predefined policies of the users are used to find out conflicts for contents and a middle ground is found with a negotiation mechanism. Such and Criado [93] extend that work by modeling user behavior and using a software mediator where some negotiation actions are made without direct user input. Kekulluoglu *et al.* [50] use multi-agent negotiation as well but apply a more comprehensive negotiation protocol. They further take into account incentives for agents. Kokciyan *et al.* [54] use argumentation to resolve privacy disputes. In that work, OSN users were represented by software agents, where the agents have access to domain knowledge and infer semantic rules. With argumentation mechanism, agents can attack other agents' assumptions to convince the others to accept its users' privacy requirements. While this is successful, it requires agents to be able to reason semantically, which may not be possible in various environments.

Mester *et al.* [63] propose four desirable properties for a privacy management system, namely automation, fairness, concealment of privacy concerns and protection before exposure. Automation property is the capability of the system to perform without human interference. Fairness property depicts the system's ability to provide

Table 2.4: *Comparison of Desirable Properties*

Properties	ACM	Negotiation	Argumentation	PCTA	PANO
Automation	✗	✓	✓	✗	✓
Fairness	✓	✗	✓	✗	✓
Concealment	✗	✗	✗	✓	✓
Protection before Exposure	✓	✓	✓	✓	✓
Easy-to-Compute	✓	✗	✗	✓	✓
Robust to Cold Start	✗	✓	✓	✗	✗
Dynamic Privacy	✗	✓	✓	✗	✗

equality to its members. Concealment of privacy concerns is about the coverage of the members' privacy policies; the more agents can hide their policies from the others, the more the property is satisfied. Protection before exposure is related to the privacy management method's application before the related entity is published in the system, causing possible violations. We add three more desirable properties for evaluating privacy management approaches, namely being *easy-to-compute*, *robustness against the cold start problem* and *dynamic privacy requirements*. Being easy-to-compute is an aspect that is important for domains where many privacy action decisions are dealt with at the same time, and when agents have limited computational power. IoT and OSNs both fit into this description, and employing easy-to-compute methods becomes essential. Being easy-to-compute does not just cover the computational complexity, but also considers time limitations under which a decision has to be taken. Having few iterations that can finalize privacy decisions in short time is advantageous. For example, a negotiation method can simply have a majority voting approach, but if it is done with several iterations with an agent hierarchy, it can still be counted as not an easy-to-compute model. The robustness to cold start property considers the success of the systems about dealing with the cold start problem, where ideally a new agent that enters the system should be able to make use of the system as well as others already present. Finally, accommodating dynamic privacy means that an agent should be able to adapt when its user's privacy requirements or the behavior of other agents change over time.

We compare our method with other privacy management systems in terms of the defined properties. We select generic approaches that enable collaboration rather than specific work in the literature. Table 2.4 shows the comparison between methods. The approaches that we compared PANO with are namely access control models (ACM), negotiation, argumentation and pure Clarke-Tax auctions (PCTA) without our modifications. In comparisons with generic approaches, we indicate that they satisfy a property even if some works include the properties and some others not.

The access control models (ACM) for privacy mechanisms depend on access rules, mostly defined as policies [38, 46]. This requires input, either by a higher authority or by users to define rules for whom to access which properties. This approach also requires revealing own policies to others, since an overseer mechanism is a must to apply and manage access policies. Therefore, ACM mechanisms do not satisfy automation, concealment, robustness to cold start and dynamic privacy. The negotiation methods for privacy management can be designed to be fully automated, even

though some methods require human intervention. Since agents have to expose their privacy requirements for negotiation, they cannot offer the concealment of privacy concerns property. Negotiation protects the agents before exposure, but it might require multiple iterations over an algorithm to finalize a privacy policy, thus easy-to-compute property is not satisfied. Negotiation approaches also might not force fair solutions, since the fairness mostly relies on the compromising of the negotiating agents. Argumentation approaches are similar to negotiation in every other property, but in addition, they also provide fairness, since the agents are given the opportunity to attack others' arguments over multiple iterations to decide the final policy, hence defending themselves against unfair arguments. A pure Clarke-Tax approach without the modifications we proposed can still satisfy the concealment, protection before exposure and being easy-to-compute properties, but it does not support automation at an applicable level, and it would require policy definition and modification by human input. Therefore, it cannot satisfy the robustness against cold start property, either. As mentioned in Section 2.2.2, even though Clarke-Tax is a mechanism that supports truthfulness, if the system is open to abuse, it fails to provide fairness. Our measures to prevent abuses to present a fair mechanism is explained in Section 2.3 and the ability of PANO to reach automation via agents enable PANO mechanism to satisfy the automation and fairness properties that a pure Clarke-Tax privacy management approach cannot. PANO conceals the privacy policies from the other participants and the bids in the auctions are blind where each participant can only know their own bid and the outcome of the auction. Therefore, PANO satisfies the concealment property. Protection before exposure is also achieved with PANO since the privacy decision is made before a content could become public, if the outcome of the auction is to share. PANO is also easy-to-compute, where it offers one-shot auctions where agents can simply pick from a set of privacy actions and bid within well-defined boundaries according to the privacy policies of the users they represent. However, the agents described in this chapter do not learn privacy requirements or behaviour of the OSN users. Therefore, they cannot adapt themselves according to varying privacy conflicts or assist OSN users in a way that the agents can work without relying on user input. Due to these drawbacks, PANO agents without a learning component still do not satisfy the robustness to cold start and dynamic privacy properties. We tackle these drawbacks in the following chapter, where we develop personal assistant agents that can learn user behaviour and autonomously act on behalf of OSN users.

As an ongoing research, PANO can be expanded to used in different domains that have the notion of collaborative privacy, such as Internet of Things (IoT) or cloud services that enable collaboration upon editing documents. One of the future directions could be to apply PANO to such domains and evaluate if it can scale up to different domain specific challenges. Another future direction can be to develop a software tool that can be used as a testbed for resolving privacy conflicts, where PANO can be compared with other approaches such as rule-based policies, negotiation or argumentation. Since the privacy domain lacks such widely accepted tools for performance evaluations, current research efforts usually tend to develop their own experiments and not being able to do comparisons on a large scale. Designing and conducting a user study to explore real-life use cases of PANO can also help to investigate the applicability of such collaborative privacy mechanisms in widely used OSNs.



3

Learning Agents for Supporting Users in Preserving Privacy

ABSTRACT

When a piece of information pertains to multiple individuals, how to share this information and with whom becomes challenging. The problem becomes more difficult when the individuals that 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.

We have proposed PANO as a collaborative privacy resolution mechanism in the previous chapter, which enables all users in a system to participate in privacy decisions for co-owned content. However, the users can lack the level of knowledge or the motivation to participate in such decisions, which brings out the need of automated agents that can act on behalf of the users they represent.

This chapter presents a personal assistant to help end-users with managing the privacy of their co-owned content. When this content 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 then produces bids accordingly. Our proposed personal assistant is capable of assisting users with different knowledge and motivation levels 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.

3.1 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 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. In online social networks (OSNs), users can share content such as images or videos that include other users, who are many times able to tag themselves or others, comment on it, and even reshare it with others. Since the pieces of content in question can contain information of more than a single user, the dissemination of content can affect the privacy of multiple users. Therefore, the content can be defined as *co-owned* [46, 95], and a collaborative privacy decision becomes essential.

In Chapter 2, we have focused on how to provide a collaborative mechanism that can be used by all co-owners of a content, with which they can describe their privacy requirements and reach a decision where privacy violations can be prevented. We have proposed a mechanism called PANO, which is built upon Clarke-Tax mechanism and lets all co-owners to enter auctions and bid on their preferences to reach privacy decisions. As shown in Chapter 2, PANO ensures that the users cannot abuse the system, and each co-owner has a similar amount of effect in the privacy decisions in the long run, where *oversharing* and *undersharing* are reduced from the native Clarke-Tax approach [107, 112].

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 on user participation in these mechanisms. First, many existing work show that users themselves do not usually know their privacy constraints, let alone evaluating the importance of contextual properties for privacy [1, 33]. 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 [47]. 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 context of privacy before, mostly to enable agents to classify whether a user would consider the content in question private or not [35, 92]. 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 for the user, what the user's valuation of the content initially was and whether the user conforms to the group 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 chapter 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, which is robust and can thus accommodate a large number of decisions to be taken. PANOLA can make use of user input on previous privacy decisions as an initial point to bid but then learns to adjust its bidding strategy over time. 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 [55]. 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 a 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 being violated.

The rest of this chapter is organized as follows: Section 3.2 explains the background on privacy personas and the common personas in OSNs. It also introduces our running example. Section 3.3 describes how PANOLA works, with a focus on its learning. Section 3.4 explains our experimental setup and answers our research questions through multiagent simulations. Finally, Section 3.5 discusses our work in relation to existing methods in the literature and gives pointers for future work.

3.2 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*. The technical aspects of PANO are explained in the previous chapter, therefore here we omit the information related to PANO and refer the readers to Chapter 2 for the details.

3.2.1 Privacy Personas

Online social networks are widely used throughout the world, with billions of users with varying understanding 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* [55]. Dupree *et al.* [33] 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 her privacy choices in the system. For example, changing privacy setting for some content or spending time to check who can access content would rank high in this dimension. As depicted in Figure 3.1, Dupree *et al.* [33] 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. Ideally, the agents that we develop should help individuals with lower knowledge and motivation as much as those with high knowledge and motivation. Only then, the system can establish *equity* and treat all users fairly [51]. 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

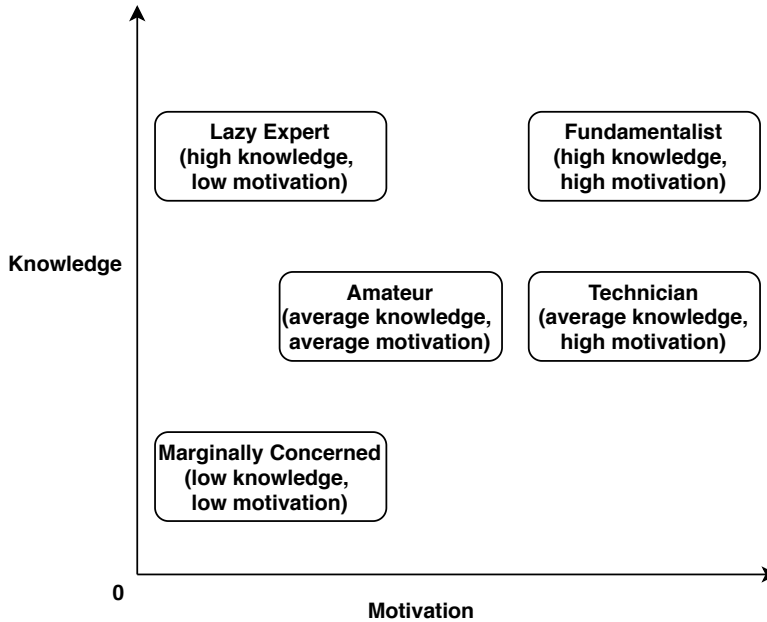


Figure 3.1: Knowledge and motivation dimensions for privacy personas, according to Dupree et al. [33].

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.

3.2.2 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 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.

Table 3.1: Example of a decision by various privacy personas

Name	Privacy Persona	Actual Pref.	Expressed Pref.
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	-

Once the PANO auction commences, each co-owner would need to assess what privacy outcome would be more fitting to their privacy understanding and place a bid for that outcome in hopes of affecting the final decision in his/her favor. Table 3.1 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 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 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.

3.3 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., tagged users), as well as the contextual properties of it (e.g., time of day or location) and then provides 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 3.2 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 rectangle 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). Afterwards, 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

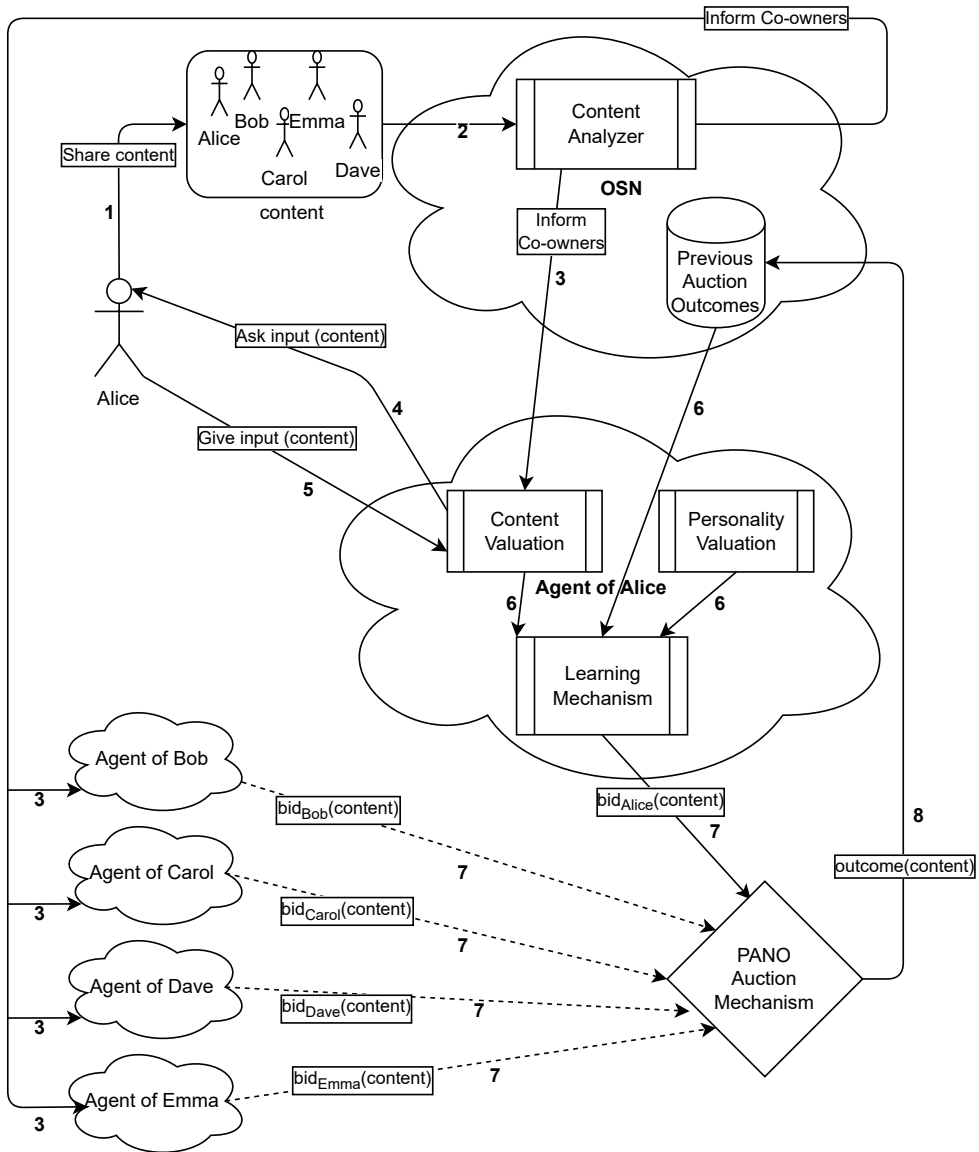


Figure 3.2: Flow diagram depicting how a PANOLA agent learns and performs bidding outcomes.

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 this 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.

3.3.1 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 straight forward way. Following this, in this chapter, 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 afterwards. 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 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, which enables agents to maximize their rewards from their actions [97]. 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.

3.3.2 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 from. 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 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 $R = \{r_1, \dots, r_n\}$.

Definition 4 For each auction outcome $t \in \{\text{share}, \text{notshare}\}$, 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; 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 decrease the computations. One solution could be to hold the PANO auctions with a small M so that the number of ranges decrease. Another solution could be to enforce a minimum length g on the range $[k, l]$ in Definition 3 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 or with the other agents. Therefore, agents can update their utility values without




Figure 3.3: Two ranges between minimum (m) and maximum (M) bidding boundaries and the initial bidding evaluation of content a for outcome t

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 units 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 7 Figure 3.3 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, which can be interpreted that the utility for r_1 was bigger than the utility for r_2 when the bid was placed.

3.3.3 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 [15, 31, 98], 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_{C_t}) 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. While conforming to others' choices can result in a user to not have her individual preferences considered for a privacy decision, it would enable a user to respect others' privacy when she does not expect a privacy violation for herself. Moreover, a user's willingness to conform to others' choices can also affect her status in a society positively, which can result in others to conform to the same users choices in the future when she needs to protect her privacy. 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 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 3.2 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 3.1:

$$C = 1 - e^{-p/S} \quad (3.1)$$

We adopt Equation 3.1 as a variation of the aging curve formula from the literature [117]. Our equation differs in the way that while the original aging curve value starts

Table 3.2: Values for utility calculations

Name Abbreviation Range	Short Description	Equation/Function
Value of Content V_{Ct} [0, 0.5]	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}
Value of Conformism V_{Cf} [0, 0.5]	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}
Confidence C_t [0, 1]	Used for defining how confident the agent is about user input for outcome t	$C_t = C \times c_t/p$ ($C = 1 - e^{-p/S}$)
Effectiveness E [0, 1]	Calculates effectiveness of a range	$E(r) = 1 - ((0.5 - V_{Ct}) \times \frac{b_{t,a}}{M} + V_{Cf} \times \frac{Tax}{M})$

from 1 and decreases over time, the confidence value in our equation starts from 0 and increases over time. According to Equation 3.1, 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 with a lower S value, the agents prematurely become confident about user input, which in some cases result 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 the 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).

$$C_t = C \times c_t/p \quad (3.2)$$

When the outcome of an action is the same with 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 3.3.

$$E(r) = 1 - ((0.5 - V_{Ct}) \times b_{t,a}/M + V_{Cf} \times Tax/M) \quad (3.3)$$

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 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. Depending on these values, effectiveness can serve as either a reward or a punishment for a range. With a lower value for content, the amount paid from the range becomes more important and lower ranges are rewarded more while bids from higher ranges are punished. In the same manner, a higher value for conformism places more importance in the paid tax, therefore becoming rewarding for the bids from the ranges that are lower in tax, and punishing for the bids from the ranges that result in higher paid tax.

3.3.4 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 3.3) 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 3.2).

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

$$Utility(r, t) = \left(\frac{\sum_{i=1}^n E_i(r)}{n} + (1 - |C_t - m\hat{e}an(r)|) \right) / 2 \quad (3.4)$$

According to Formula 3.4, 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 ($m\hat{e}an(r)$) 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 as well as the user input are considered. We give an example below to demonstrate how the confidence value would affect the selected range.

Table 3.3: *Utility update examples for Example 2*

Ex.	Effectiv.	Distance of range from conf.
#1	$r_1 = 0.3$	$r_1 : 1 - 0.1 - 0.15 = 0.95$
	$r_2 = 0.9$	$r_2 : 1 - 0.1 - 0.85 = 0.25$
#2	$r_1 = 0.3$	$r_1 : 1 - 0.5 - 0.15 = 0.65$
	$r_2 = 0.9$	$r_2 : 1 - 0.5 - 0.85 = 0.65$
Utility		
$Utility(r_1, notShare) = (0.3 + 0.95)/2 = 0.625$		
$Utility(r_2, notShare) = (0.9 + 0.25)/2 = 0.575$		
$Utility(r_1, notShare) = (0.3 + 0.65)/2 = 0.475$		
$Utility(r_2, notShare) = (0.9 + 0.65)/2 = 0.775$		

Example 8 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 on Table 3.3. Using Equation 3.2 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 $m\hat{e}an_{r_1} = 0.15$ ($3/20$) and $m\hat{e}an_{r_2}$ 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.

3.3.5 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 3.2, agents interact with the OSN to make privacy decisions, since the PANO auctions are governed by the OSN.

Algorithm 2: 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( $\langle decision, tax \rangle$ )
17  $E \leftarrow E \cup E(best)$  //Equation 3

```

Algorithm 2 explains the steps the PANOLA agents take throughout the decision making process. When a new piece of content is introduced to the OSN which 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 3.2), 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 3.4 (line 8) as explained in Section 3.3.4. 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 which 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 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 which the agent is able to bid from with its owned budget. Afterwards, 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 3.3) to the set of all previous effectiveness values (E) in line 17.

3.3.6 Running Example with PANOLA

We now walk through our running example from Section 3.2.2, 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 3.4 and Table 3.5 show the information about the users and their personas, their prior input and for which output they placed their bids for both examples, respectively.

Table 3.4: *Example of a decision with no prior knowledge*

Name	Privacy Persona	Input	Actual Pref.	Confidence	Bid
Alice	Lazy Expert	-	Not Share	0	-
Bob	Margin. Conc.	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	-

In the first iteration of the example in Section 3.2.2, the agents have no prior information. Bob, Carol and Dave are the users to give input as seen in Table 3.4. Recall that Bob has the *Marginally 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.

Table 3.5: *Example of a decision after 20 previous auctions*

Name	Priv. Persona	Input	Act. Pref.	Confidence	Bid
Alice	Lazy Expert	0 Share 7 Not Share	Not Share	$C_t = 0.50$	10, $t = \text{NotShare}$
Bob	Margin. Conc.	2 Share 6 Not Share	Not Share	$C_t = 0.41$	7, $t = \text{NotShare}$
Carol	Fundament.	18 Share 0 Not Share	Share	$C_t = 0.83$	16, $t = \text{Share}$
Dave	Technician	3 Share 14 Not Share	Not Share	$C_t = 0.67$	13, $t = \text{NotShare}$
Emma	Amateur	3 Share 10 Not Share	Not Share	$C_t = 0.56$	11, $t = \text{NotShare}$

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 3.5, 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 on the confidence column of Table 3.5. 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 advocate 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.

3.4 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 afterwards an equity of the decisions should be seen.

Research Questions: We formulate with 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 Dupree *et al.* [33]’s 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 Dupree *et al.*’s 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, 0 representing the lowest knowledge and motivation and 100 being the highest. To be in line with knowledge and motivation dimensions represented in Dupree *et al.* [33], we assign three levels for both, which can differ for each persona. Table 3.6 shows the knowledge and motivation levels for each persona, as well as their percentage in the entire community. According to Table 3.6, 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>

Table 3.6: *Percentage of all personas in our evaluations and their knowledge and motivation levels*

Privacy Persona	Prcent. 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%

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 that 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 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 which 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.

$$s = \frac{SA}{n} \quad (3.5)$$

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 which include the desired true parameter (the arithmetic mean in our case) with a confidence level. Confidence intervals suit our experiment well [32], since they demonstrate the probability of the deviation from the results. In order to evaluate the confidence intervals, we employ the following formula:

$$\bar{X} \pm Z \times \frac{\sigma}{\sqrt{n}} \quad (3.6)$$

where \bar{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 [32]; hence, the Z value is assigned as 1.960.

3.4.1 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 3.3.1.

We run two experiments to evaluate the success of each privacy persona represented in the simulation by PANOLA agents or non-learning agents. For each setup, we execute 50 runs for each persona, each with 11000 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 3.6. In addition to five personas, we run the same setup with a *random agent* to compare the performance of PANOLA agents with this base case. We use the first 1000 items of content for the learning phase of PANOLA agents and therefore do not measure the success for those. The auctions for the remaining 10000 items of content are used for the test set; thus, they are a part of the evaluations.

Figure 3.4 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 3.4(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


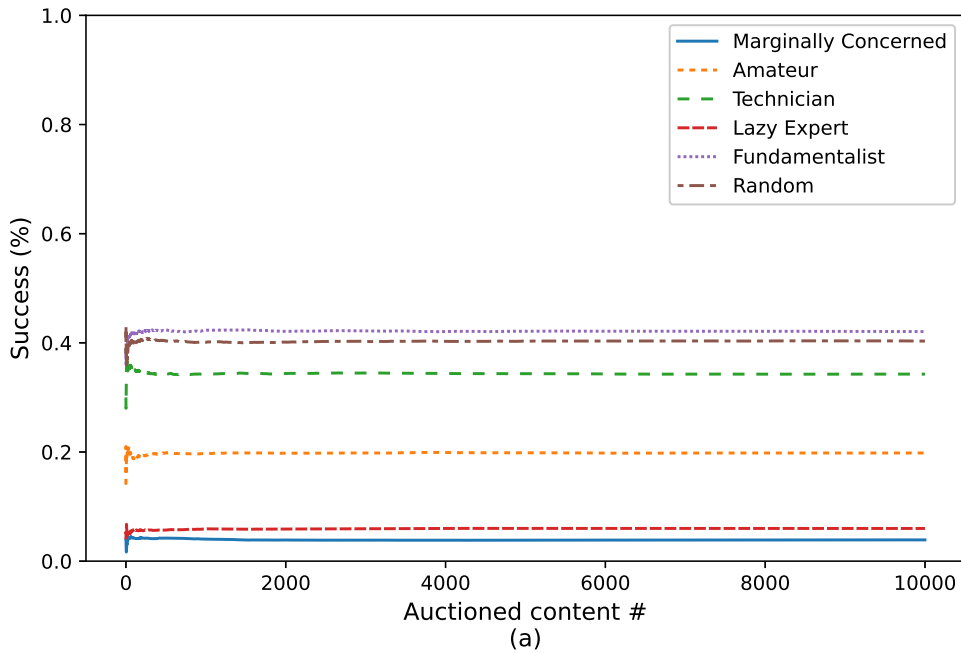


Figure 3.4: Success of (a) non-learning agents against each other (b) PANOLA agents and random agent against non-learning agents.

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, she also gets taxed, which would leave 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 3.4(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 11000 auctions. Since the agents already go through a learning phase for the first 1000 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 6000 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 the 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 3.4, we 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 3.7 shows our results of statistical significance with confidence intervals for the final results of this setup. According to Table 3.7, we can see that the experiments from Figure 3.4, 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.

Table 3.7: Statistical significance analysis of experiments depicted in Figure 3.4, after 50 runs per experiment

Agent Type	Learning Type	Mean of Success	Confid. 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

3.4.2 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, 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 11000 items of content, 1000 for training and 10000 for test of success, and we perform 50 runs with the same setup for each persona.

Figure 3.5 shows the results of this experiment. The trend of success is similar to the results from 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 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 which 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 10000 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 3.8 depicts the results of statistical significance analysis with confidence

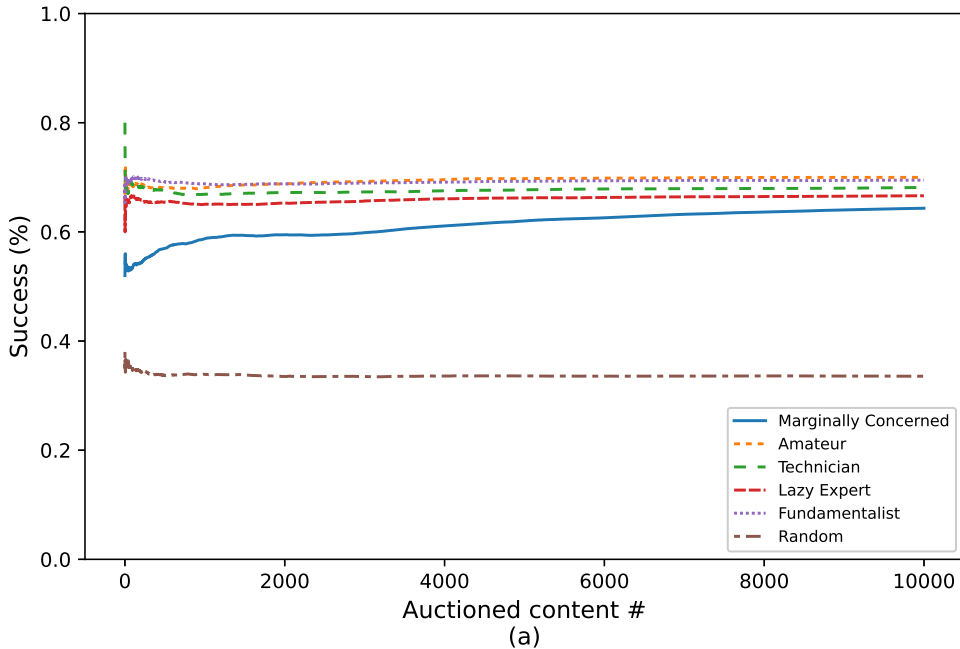


Figure 3.5: Success of PANOLA agents against other PANOLA agents with varying privacy personas.

intervals for this experiment. When all the agents are PANOLA agents, the margin of error slightly increases than 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.

Table 3.8: Statistical significance analysis of experiments depicted in Figure 3.5, 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

3.4.3 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 3.4.1. 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, thus are more active in the auctions which 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 differently 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 1000 pieces of content and over 50 runs we will demonstrate how quick they converge to find out the possible winning bids.

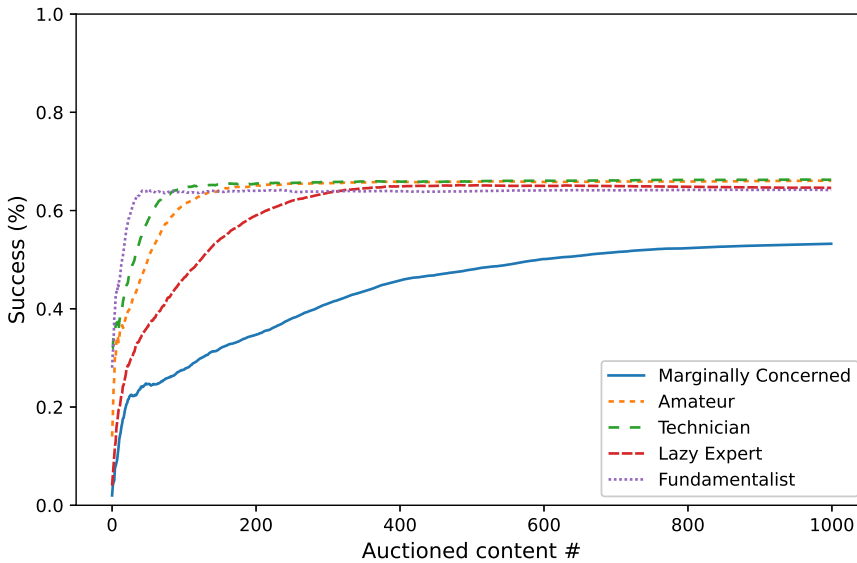


Figure 3.6: Learning progress of PANOLA agents with each privacy persona over 1000 pieces of content.

According to Figure 3.6, we can see that the trends match with the outcome of Figure 3.5, which shows the results after the learning phase. We can also conclude that both knowledge and motivation levels affect the speed of reaching a stable point. Since

Fundamentalists have the highest knowledge and motivation levels, PANOLA agents which 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 afterwards. 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 1000 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 to become more confident about how to perform in auctions.

Table 3.9 depicts the results of statistical significance analysis with confidence intervals for the learning evaluations. For 1000 items of content that is shown in Figure 3.6, 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.

Table 3.9: Statistical significance analysis of experiments depicted in Figure 3.6, 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

3.4.4 Employing Values for Individualism or Conformism

In Section 3.3.1, we have explained two parameters in efficiency calculation for bid ranges, namely *Value of Content* (V_{C_t}) and *Value of Conformism* (V_{C_f}). As explained in Section 3.3.3, a higher valuation of V_{C_t} means that PANOLA agents do not mind the amount of bids they place when content is valuable for them, since a higher V_{C_t} enables the agents to consider that winning an auction is more important than spending budget. On the contrary, when the V_{C_t} is lower, agents try to minimize their bids as much as possible in the learning process and therefore the agents who set V_{C_t} low might lose some auctions while trying to find the winning bid. With V_{C_f}

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. As discussed in Section 3.3.3, while conforming to others' privacy choices can result in an agent to not achieve all its privacy requirements, it would give the agent the ability to respect the privacy of others, especially when the outcome would not explicitly violate its privacy.

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} a 0. With these two valuations, we investigate two different number of co-owners to see the effect of the level of opposition. For the first co-owners setup, we assign 3 co-owners who have privacy preference opposing the PANOLA agent and only one with the same privacy preference. For the second setup there are 2 non-learning agents against the PANOLA agent instead of 3, 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 10000 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 3.4.1, 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 3.7), 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.

$$SS = \frac{n - d}{n} \quad (3.7)$$

Figure 3.7 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 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,

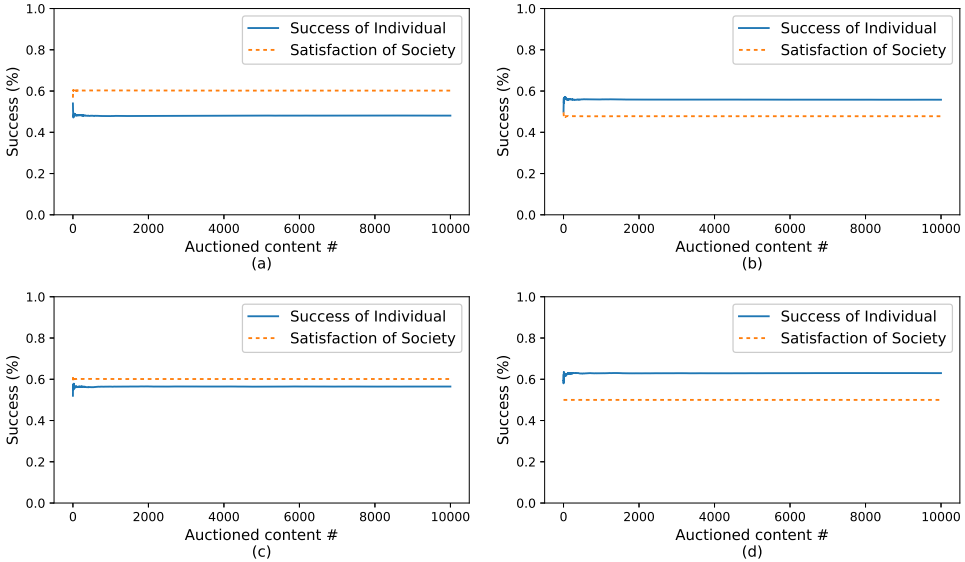


Figure 3.7: Success of PANOLA agent and satisfaction of society in (a) 1 agent on the side of a conformist PANOLA agent against 3 agents. (b) 1 agent on the side of an individualistic PANOLA agent against 3 agents. (c) 1 agent on the side of a conformist PANOLA agent against 2 agents. (d) 1 agent on the side of an individualistic PANOLA agent against 2 agents.

they 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.

3.5 Discussion

We have explained how PANOLA agents can assist OSN users in order to protect their privacy for collaborative decisions. Then, we have 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 3.4). 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 3.5 that, when all users employ PANOLA, the success rate of the users—almost independent of the persona—converge 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 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 her preferences. Similarly, the agent could also get input from the user about her satisfaction with the previous interactions to further improve the learning process. In this chapter, 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 [89] 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.

Referring back to the desirable properties that were discussed in the previous chapter, we have stated that PANO satisfies most of them, with the exception of being robust to cold start problems or having a dynamic approach to resolving privacy disputes. PANOLA agents keep the abilities of PANO agents, therefore satisfy all the properties that PANO were able to do. In addition, with the introduction of PANOLA agents which can assist OSN users, we would be able to resolve cold start problems, since PANOLA can enable successful collaborative privacy decisions even when the provided feedback is very limited. For the dynamic privacy property, we have shown that PANOLA agents adapt their bids when the other parties in the auctions change their behavior, while maintaining a steady success rate to support the privacy requirements of the OSN users they assist. We show the updated desirable property comparison in Table 3.10 with PANOLA agents, compared with access control models (ACM), negotiation approaches, argumentation approaches, pure Clarke-Tax auctions (PCTA) and PANO without learning agents.

3.5.1 Related Work

With the Internet becoming accessible to most people in the world around the early 2000s, it enabled easy sharing of and access to private information over the web, which brought privacy to attention in the systems that make use of the internet. Langheinrich [57] is one of the first studies to investigate the open issues for privacy-respecting approaches for ubiquitous computing, which can be defined as the use of computing in any device, in any place and in any format. Spiekermann and Cranor [87] and Gürses *et al.* [44] study the grounds of engineering privacy, explaining how informa-

Table 3.10: *Comparison of Desirable Properties*

Properties	ACM	Nego.	Argum.	PCTA	PANO	PANOLA
Automation	✗	✓	✓	✗	✓	✓
Fairness	✓	✗	✓	✗	✓	✓
Concealment	✗	✗	✗	✓	✓	✓
Protection before Exposure	✓	✓	✓	✓	✓	✓
Easy-to-Compute	✓	✗	✗	✓	✓	✓
Robust to Cold Start	✗	✓	✓	✗	✗	✓
Dynamic Privacy	✗	✓	✓	✗	✗	✓

tion related domains can be designed to employ privacy-preserving methods. Paci *et al.* [74] 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.* [12] show the challenges of preserving privacy over decentralized OSNs and provides a review of previous work done for overcoming these challenges. Bertino and Ferrari *et al.* [17] 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 [112] a 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 [110]. 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 [94] 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 chapter 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 [38] 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.* [46] is another work which 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 [107, 112]. 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.* [52] 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 [95] 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 [93], and the same approach was extended with modeling and learning user behavior, as well as implementing a software mediator which manages the negotiation process without the need of a human interruption. Another multi-agent negotiation model was introduced in Kekulluoglu *et al.* [50], which includes a comprehensive negotiation protocol to be used by the agents. In addition, incentives of the agents are considered in the approach.

Kökcüyan *et al.* [54] 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. How-

ever, 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 are gaining traction. Dechesne *et al.* [30] uses a case example from law to model the effect of culture in human societies, in terms of values and the acceptance of and compliance with norms. Calikli *et al.* [19] 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.* [3] study norm emergence factoring in the context of the agents, taking the sanctions into account. In another work, Ajmeri *et al.* [4] provide a framework where agents aggregate the value preferences of the users and choose ethically appropriate actions for social contexts. Ulusoy and Yolum [111] propose a norm-based access-control mechanism for collaborative privacy decisions, which considers both personal and social norms in decision making. Mosca *et al.* [66] 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.* [23] 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. Akata *et al.* [5] describes a research agenda about *hybrid intelligence*, which considers values such as adaptability, responsibility and explainability to form a bridge between human and machine intelligence in order to achieve goals that are unobtainable by both parties when are not combined.

The use of machine learning for privacy is gaining momentum, and the research area is still open for further improvement. Fogues *et al.* [36] provide an agent-based approach which 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.* [114] 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.* [91] 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.* [120] 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.* [92] aim to learn privacy features of image content, with a novel approach of employing sentiment in the context of image classification. Fang and LeFevre [35] 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.* [8] present another privacy policy learning approach, which is based on ReBAC [38]. 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.

3.5.2 Future Work

In light of this work, some interesting research directions open up. Modelling 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 will be focusing on Chapter 5. Use of social norms could be beneficial to create learning agents according to their normative behavior that can also benefit 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 behaviours 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 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.



4

Privacy Values in Collaborative Systems

ABSTRACT

Employing software agents as personal assistants is beneficial for collaborative privacy decisions, since they can significantly reduce the need for user interaction for privacy and can help providing privacy preserving decisions for users with different knowledge and motivation levels. The agents can learn how to act on behalf of the users with user feedback, or by making use of previous decisions made in the system.

While privacy assistants are useful for privacy preservation of the entire community, their main goal would be to protect privacy of the users they represent, even when those users have varying values and would be willing to act differently.

As in every society, humans value different aspects in their decisions. While some prefer individualism, others can behave in an altruistic manner, sacrificing their own goals for the greater good. A fraction of humans prefer to stay conservative and stick to their initial decisions, while some are open to changes and search different alternatives.

Since online social networks represent a society, we would expect these values to be present in their users. Therefore, investigating how different learning strategies that are affected by these values can provide some insights for how human values should be incorporated in software agents that can assist the users.

In this chapter, we evaluate various scenarios in simulations where software agents represent different human values. The human values are adopted from the scientific literature and evaluated when they are on the opposite scales of a comparable value. We present the results along with insights on how they can be employed for different agents in different domains.

4.1 Introduction

In Chapter 3, we have explained and evaluated PANOLA, which assists online social network users for collaborative privacy decisions, and help them to preserve their privacy even without the need of user input. While PANOLA can provide an efficient way to reach privacy related decisions regardless of the knowledge level of the users, or their motivation to participate in such tasks, the main aim stays as to preserve privacy. In human societies, people can value different aspects for decision making, thus might be willing to act differently than others. For the privacy related decisions, some might value sharing content to such an extent that giving up on their privacy is preferable. This can create a trade-off, where human values would play a role in how a user in an OSN would like to behave. The trade-off between sharing content and preserving privacy has many interesting directions that pertain to human behavior, especially as to what one sees important in her interactions. For example, some people would rather help others preserve privacy than share content themselves, while others would fight to share. Or, while some individuals change their privacy expectations over time, some others would never give up their initial expectations. These characteristics are the *human values* that shape privacy interactions. Ideally, a software agent that assist a user should be able to capture its user's values in relation to privacy and adapt its learning accordingly. In this chapter, we investigate how the values can affect the learning strategies and aim to give pointers for fine tuning of PANOLA for different domains.

In order to represent values for privacy, we start with Schwartz' model of human values [81]. Schwartz establishes a model based on several works of theorists, and defines six features about the nature of values. These features are listed as below:

- Values are beliefs, which infuse feelings for the people who has them when they are realized.
- Values are linked with desirable goals, which motivates actions to reach them.
- Values go beyond specific situations and can induce similar actions in varying cases.
- Values guide the actions of the people when a decision is to be made.
- In case of conflicting values, people tend to give more importance to some values over others and take actions in line with the most important one.
- Each taken action can be in relation to more than one value, and the relative importance of multiple values guide the actions taken.

In light of Schwartz' features for values, we can say that the actions taken for privacy decisions are also aligned with the same features. For example, if a person values social appreciation, she might be willing to take privacy actions mostly inclined with sharing content. In an opposite case, a person who values individual privacy might be willing to sway others' decisions to not share a group picture. Schwartz defines ten motivational types of values, and places them into four main dimensions, as shown




Figure 4.1: Schwartz' dimensions for human values [81].

in Figure 4.1. These four dimensions are *Openness to change*, *Conservation*, *Self-enhancement* and *Self-transcendence*, where the dimensions reflect a trade-off in pairs but the pairs themselves are orthogonal. Openness to change reflects the willingness of a human to change her behavior over time, whereas the conservation reflects stubbornness. In context of privacy, we interpret self-enhancement as a person's willingness to enforce her own constraints and self-transcendence as her willingness to accommodate other people's choices.

The alignment of Schwartz value definitions with multiparty privacy approaches has already been considered for some applications in the literature. Mosca *et al.* [69] adopts some of the Schwartz values for resolving multiparty privacy conflicts where users have an order of importance over these values in their decisions. Mosca *et al.* also used Schwartz values for value driven explainable agents in collective privacy [66, 70]. Mosca and Such [68] extended this work to develop an explainable personal assistant which models user privacy behavior in accordance with Schwartz values.

In terms with Schwartz's human value definitions, defining and employing various agent characteristics for PANOLA can be beneficial to model the privacy behavior of users, since users will behave differently according to their values, as expected. In

this chapter, we investigate four types of agent characteristics that might differ from agent to agent, which are explained below. In order to focus on the effect of values, we exclude privacy personas for the OSN users, and assume that each user is driven by their values with an equal level of knowledge and motivation.

Openness to Change vs. Conservation: To capture this value in the agent, we enable the agent to search their bidding ranges to decide on a bid using different distributions. *Random* distribution enables agents to search the entire range freely, thus, open to change, whereas *Gaussian* distribution guides the agent to bid closer to the mean of the ranges to get conservative bids. We call the agent that uses a Random distribution as *Adventurous* while the agent with a Gaussian distribution as *Conservative*.

Self-enhancement vs. Self-transcendence: To capture this value in the agent, we benefit from the use of budget. Agents can pick a bid within the selected range according to their budget balance. An agent might want to spend more when it considers a content as important and have sufficient budget or prefer to spend less in the opposite case. We introduce three types of agents for this characteristic within the scope of this work, namely *benevolent*, *pragmatic* and *greedy* agents. A *benevolent* agent tries to spend less than its evaluated bid value, and allowing the rest of the society have a bigger influence on privacy decisions (i.e. *Self-transcendence*) while a *greedy* agent aims the opposite and tries to reach selfish decisions (i.e. *Self-enhancement*). *Pragmatic* agents always try to bid closer to their actual bid evaluations, and aims a balance between *Self-enhancement* and *Self-transcendence* dimensions of Schwartz values.

In OSNs, users might value content in different contexts differently. Some content might be more important, where users would want the privacy decision to be in their favor, while for others they might not try to impose their decisions when they deem the content not important. In our context, we use importance to define how much agents should be involved in auctions. If an agent enters in an auction and has related policies with high importance values; it indicates that the agent should try to bid a significant amount of currency for its preferred privacy actions. On the contrary, if an agent has some related policies for an auction with close to zero importance values, it shows that the agent is not willing to spend much for this auction, since the outcome has very little significance to it. In our evaluations, we divide importance into three categories, namely *insignificant*, *ordinary*, and *significant* in relation to the importance values. When the agent find the content very private, we call it *significant*. If a given content is not considered private by the agent, we call it *insignificant*. When the agent finds the content private but does not want to bid more than it earns for an auction, we call it *ordinary*.

4.2 Learning action choices

For a content and a set of possible privacy actions, an agent might have multiple actions that can satisfy its privacy requirements. For example, if an agent believes that it would not make a winning bid on its first choice against other auction participants, then bidding for secondary or tertiary privacy options might be beneficial. In order

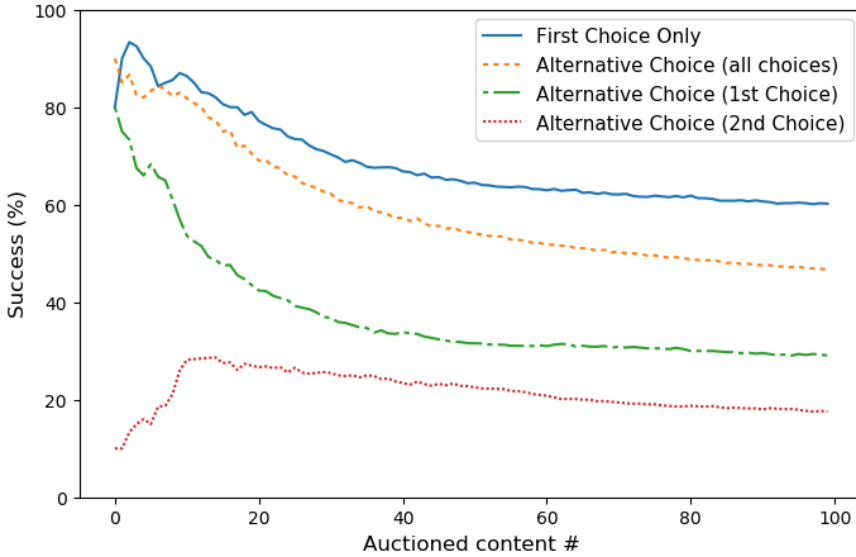


Figure 4.2: Success of First Choice Picker and Alternative Choice Picker PANOLA against non-learning agents.

to facilitate this, the agent will not only need to learn the bidding ranges for its first choice of action but also for all preferred actions. We investigate the effect of action choice characteristic, in terms of learning for only the first choice and learning for all preferred actions. We define the categories for action choice characteristic for PANOLA as *First Choice Picker* and *Alternative Choice Picker*. While the former only uses reinforcement learning as described in Chapter 3 for its top choice, the latter uses the same learning algorithm also for its alternative action choices, which we limit to a secondary action for our tests.

We compare both categories against non-learning agents, while the privacy importance of contents is set to ordinary for all agents, in a similar setup with the evaluation in Section 3.4. Figure 4.2 shows the results for privacy success percentages auctions of *First Choice Picker* and *Alternative Choice Picker* over 100 pieces of content. We also include the breakdown of Alternative Choice Picker’s success for its first and second preferred action in Figure 4.2, along with its overall success, which is the sum of these. The results show that *First Choice Picker* PANOLA ends up with ~60% successful privacy decisions against non-learning agents, while *Alternative Choice Picker* PANOLA decreases to slightly less than 50% for overall success. Recall that 100% success can never be achieved in a setup where privacy conflicts are present, since a successful outcome for a user would mean an unsuccessful outcome for another user that has a conflicting privacy preference. Moreover, *Alternative Choice Picker* can win with its secondary action choice to some extent (~20%) after learning the possible winning ranges for this option. The main reason behind this is that Alternative Choice Picker tries to win for its first and second choices when it can not win, and since these choices can conflict with each other, it can cause the agent’s second ac-

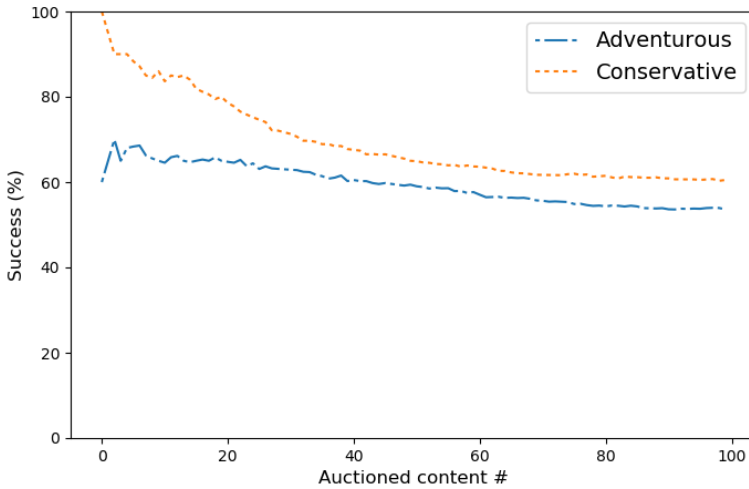


Figure 4.3: *Success of adventurous and conservative PANOLA against non-learning agents.*

tion to outbid its first for some auctions. Therefore, to win for both choices it keeps increasing its bids for, eventually causing itself to get taxed with significant amounts and run out of budget easily, where the opponents have higher chance of outbidding the agent.

Learning different actions can be beneficial for other domains where the actions are mutually exclusive, meaning that each action has different benefits and consequences. In terms of human values, these results reflect that regardless of leaning more to self-enhancement or self-transcendence, the agents should still focus on learning their primary choice for privacy while assisting OSN users in collaborative privacy decisions using PANO. According to the result of this evaluation, we conclude that learning only for the first action choice is a better strategy to employ for auctions in this privacy domain, and we will adopt first action choice strategy for the remainder of this section.

4.3 Effect of openness to change and conservation

As described in Chapter 3, PANOLA makes use of bidding ranges to learn which are the most suitable to bid from, and updating their utilities depending on the outcome. Referring to human values, PANOLA adopts a conservative strategy where it prefers less exploration within the ranges. While learning which range a bid will be given from is an important step, deciding on the actual bid within the picked range is important as well. Intuitively, the agent can pick a bid from the range based on a given distribution. Here, we implement two types of agents, namely *adventurous* and *conservative*. *Adventurous* agents bid randomly within the picked bidding range,

while *conservative* agents bid according to normal distribution in Gaussian. Adventurous agents fit into *Openness to change* dimension of Schwartz value system, since it explores all the possible bids within a range without any prior intuition about which bid could be the better one. Conservative agents are the opposite of this, and act closer to the *Conservation* dimension according to Schwartz values. This is mainly because bidding with Gaussian distribution in a range tends to get bids closer to the mean of the range, thus closer to the bid ranges are not favored, causing the agent to pick the same ranges as long as they are winning auctions.

First, we compare the adventurous and conservative PANOLA agents against non-learning agents, which always bid the same amount according to their initial bid evaluations. Both PANOLA and non-learning agents consider privacy value of a content in question as *ordinary*, thus we ensure that all agents value the incoming content at a similar importance.

Figure 4.3 shows that when all the agents consider contents equally private, the conservative PANOLA performs better than the adventurous PANOLA, against non-learning agents that do not update their privacy requirements according to their previous privacy decisions. The main reason behind this is the conservative PANOLA agent tends to pick bids closer to the selected ranges, while adventurous agent might pick the highest or the lowest bid of the selected range by chance. Since the opponent is always bidding the same amount for the privacy action it wants, learnt winning ranges by PANOLA after some privacy auctions mostly have higher means than the bids of the non-learning agents. Therefore, a conservative bid around the mean for PANOLA would have a higher chance of winning, while an adventurous bid closer to the minimum of the range might lose the auction or a bid closer to the maximum would be taxed more, causing the adventurous agent to lose more units while winning. We can conclude that in any case, learning how to bid over time is advantageous with both adventurous and conservative approaches, against non-changing privacy biddings for the desired privacy actions and that being conservative yields better privacy decisions for agents. In terms of human values, openness to change does not always bode well when it comes to privacy decisions in PANO, since it might result in paying more than a reasonable amount to reach an outcome that would be preferred by the user, or not bidding enough to have a favorable outcome.

Next, we compare the performance of the adventurous and conservative PANOLA agents against each other with differing content privacy evaluations. We use three importance privacy levels for contents, namely significant, ordinary and insignificant, as described in Section 4.1 for different scenarios to understand how different characteristics are affected by how an agent evaluates the importance of a content. The importance level enables us to also measure the different behavior when an agent values the same content differently. These importance levels are implemented as a series of scenarios when they differ for the opposing agents, where we investigate a *category x* agent characteristic with *importance level i* against *category y* agent characteristic with *importance level j* for a scenario and swap their importance levels for the next one.

In order to compare adventurous and conservative PANOLA agents for their bidding distributions, we implement five scenarios with varying content privacy importance evaluations for the agents. We investigate the success rate and total owned

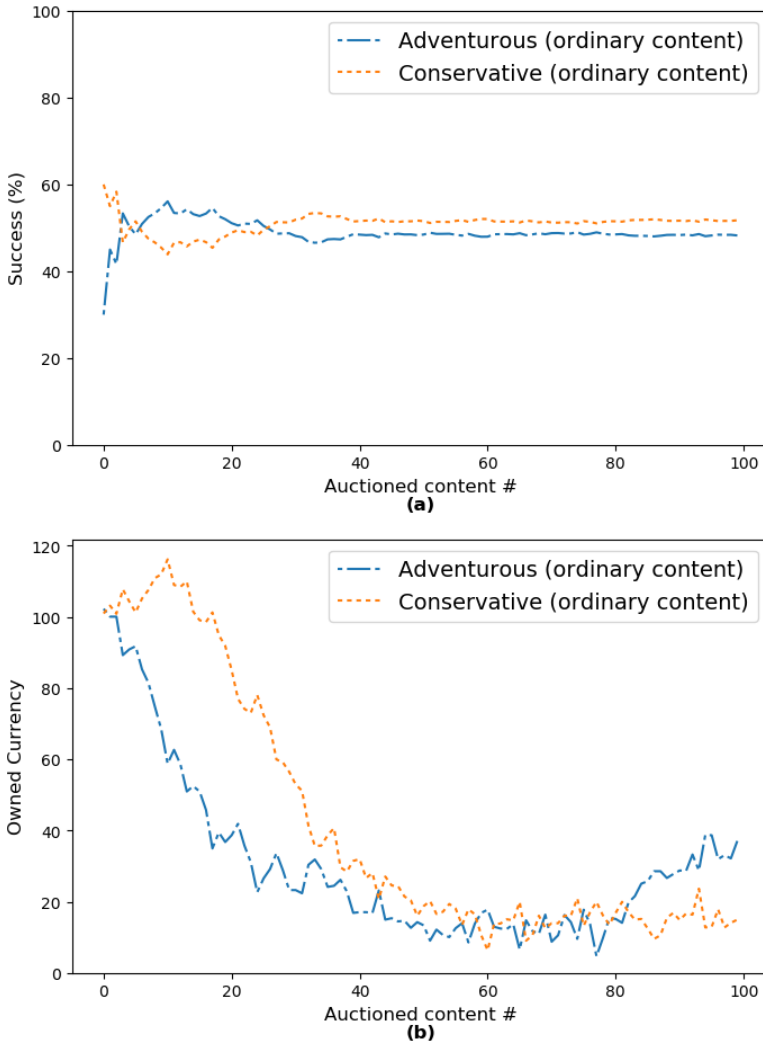


Figure 4.4: *Success (a) and Owned Budget (b) of adventurous and conservative PANOLA against each other, when both consider the content ordinary*

budget of the agents over 100 auctions.

Ordinary vs. Ordinary Content: In this scenario, both adventurous and conservative agents consider content as ordinary for privacy decision. Figure 4.4 shows the success rates and total owned budget over 100 auctions for both agents.

According to Figure 4.4, it can be seen that conservative bidding gives slightly more successful results after the agent learns the environment through some auctions. It is also more successful at the first few auctions, while spending more reasonably than the adventurous bidding with random distribution. Around the tenth auction,

adventurous agent's success passes conservative, since the adventurous agent tries to increase its bids to beat conservative, while conservative does not increase its bids since it already wins auctions. But after the next few auctions, conservative agent also adjusts its bids accordingly, and stays steadily around 4% more successful than the adventurous agent. The main reason for this difference relies on PANO mechanism; when a conservative agent outbids the adventurous, the tax amount payed tends to be a small amount, since the conservative agent sticks closer to its winning range and not reaching the maximum boundaries. In the opposite position, an adventurous agent can win by trying bids closer to the maximum boundary, but get taxed with a bigger amount which decreases its budget significantly for the next auction. According to this evaluation, it can be said that when two learning agents have the same importance evaluation for an incoming content, using a conservative approach leads to more successful bids in the long run.

Ordinary vs. Significant Content: We investigate two scenarios, where in one conservative agent considers the contents as significant and adventurous as ordinary; and in the other scenario the importance evaluations are reversed. Figure 4.5 shows the evaluation results over 100 auctions for each scenario.

Recall that a significant content means that the agent aims to bid for the privacy action it favors around 50% more than the units it earns from an auction. The results for 4.5 (a) show that conservative agent stays more successful until around 35th auction, but adventurous passes it for the next ~50 auction, only to be passed again by conservative. The reason conservative having lower success rate for around 50 auctions is that it spends more than it earns for auctions and depletes its budget quickly. However, it still learns to bid the right amount to outbid adventurous agent's bids, and adventurous agent loses auctions even though it has significantly more units than conservative agent.

The results show a different pattern than the one with reversed importance evaluations. As can be seen in The results for 4.5 (b); for about 25 auctions at the beginning, the adventurous agent's success is far greater than conservative agents' in the reversed scenario, since the random spends big amounts while picking bids from higher bidding ranges. But the conservative agent takes over afterwards and stays more successful than adventurous agent, and even saves budget doing this, while adventurous agent's owned budget depletes. The main reason behind this is that conservative agent does not vary much from its initial evaluations when it starts to win auctions, while adventurous agent can bid very high within its evaluation range and win an auction, but paying a significant amount of tax.

In light of the second and third scenarios, it can be seen that using a conservative approach against in both scenarios is more reasonable in terms of budget spending. It also performs slightly better than an adventurous agent after winning bidding ranges are learned. However, when the opponent has enough budget and also considers an incoming content significantly private, learning process of a conservative agent takes more time than an adventurous agent, causing it to lose a few auctions before winning bids are achieved. Hence, in situations where the opponent is understood to be rich on budget and winning privacy auctions, trying an adventurous strategy can perform better than learning with a conservative strategy.




Figure 4.5: Success of adventurous and conservative PANOLA against each other, for content considered ordinary by adventurous and significant by conservative (a) and for content considered significant by adventurous and ordinary by conservative (b)

Ordinary vs. Insignificant Content: We develop two complementary scenarios where we compare a conservative agent who considers incoming contents insignificant for a privacy action against an adventurous agent who considers incoming contents ordinary for one scenario and the reversed importance evaluations for the other. Similar to biddings for significant content, the bidding aim of an agent who considers a content insignificant differs from the units earned from an auction, but this time it is

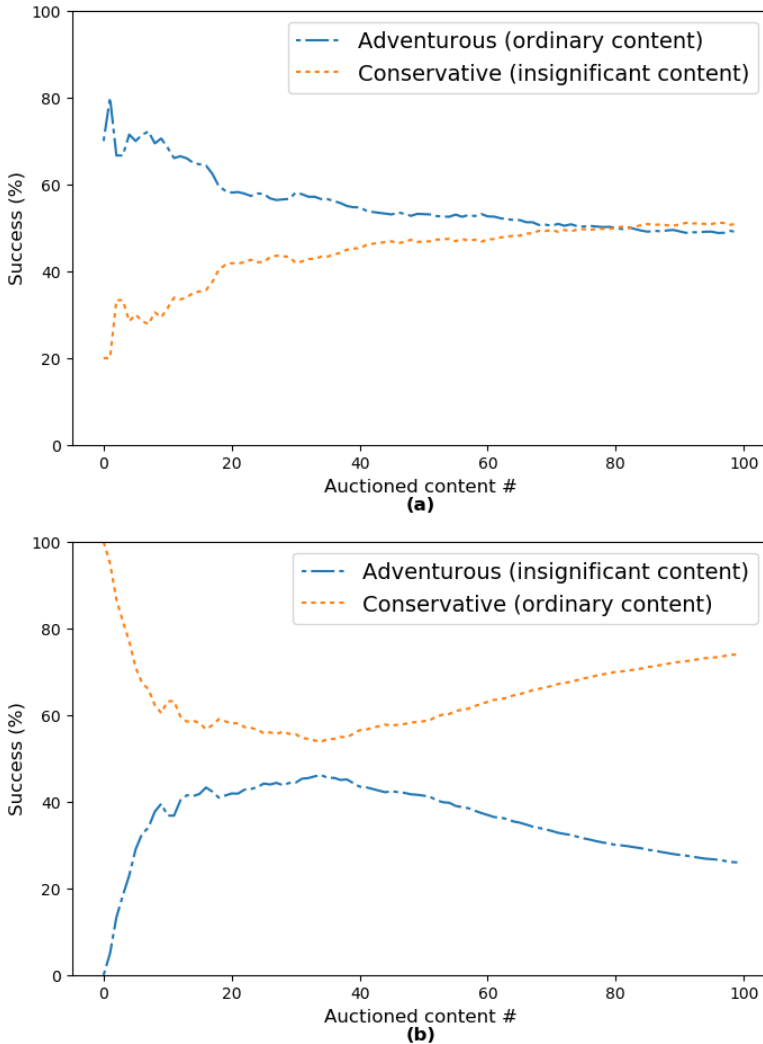


Figure 4.6: Success of adventurous and conservative PANOLA against each other, for content considered ordinary by adventurous and insignificant by conservative (a) and for content considered insignificant by adventurous and ordinary by conservative (b)

50% lower. Figure 4.6 shows the success and owned budget evaluations of this two scenarios over 100 auctions for each.

The results shown in Figure 4.6 (a) has similar patterns with Figure 4.5(b). Even though conservative agent spends less and has lower success rates, near the end of 100 auctions, it eventually passes adventurous agent at success rate, and still saves significantly higher amount of units than random.

As can be seen in Figure 4.6 (b), adventurous agent performs far worse than con-

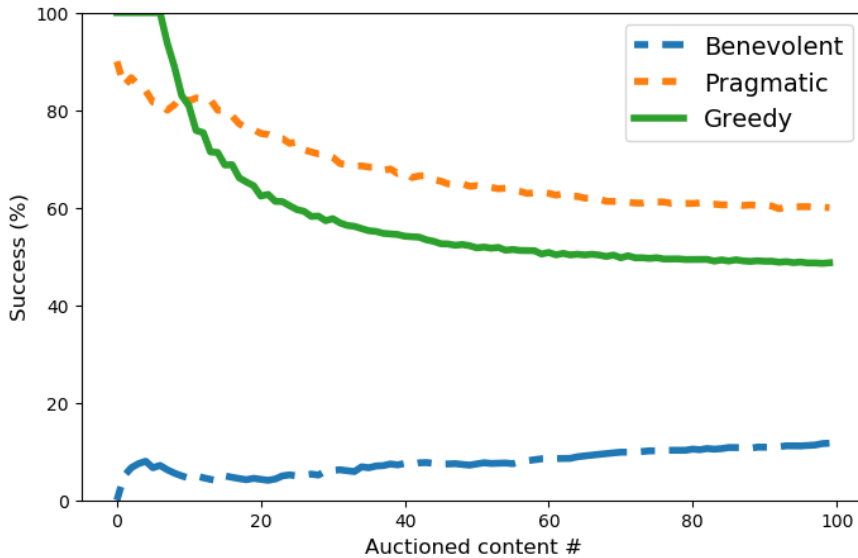


Figure 4.7: Success of benevolent, pragmatic, and greedy PANOLA against non-learning agents.

servative agent does in the same situation in this scenario. It tries to bid more and increases its success rate for a while, but eventually gives up and starts to preserve money by not bidding much in auctions. Meanwhile, conservative agent also preserves some budget while winning most of the auctions. These scenarios show that in ordinary vs. insignificant comparisons, conservative strategy is superior to adventurous one.

According to the result of all these scenarios, we can conclude that employing conservative strategy in biddings is more beneficial than the adventurous strategy in most cases. However, the learning curve of an adventurous agent while losing is steeper than the conservative one. Thus, while winning the conservative strategy is advantageous for successful privacy decisions, when the agent loses most of the bids, trying an adventurous strategy while trying to pick from higher ranges could be beneficial to find out the winning privacy bids over opponents. In terms of human values, the results show that in the long term, conservation works better for collaborative privacy decisions with PANOLA, while being open to change might help PANOLA agents to adapt quicker to others, and in short term being more successful in some cases.

4.4 Effect of self-transcendence and self-enhancement

Because of the repeated nature of the interactions for privacy decisions, agents' approach on *self-enhancement* and *self-transcendence* values becomes vital, since it directly affects how they spend their money over time. If an agent values *self-enhancement*, it tries to enforce its own privacy decisions over society, thus they

become *greedy* agents. A greedy agent spends a lot of its money to fight for the privacy of incoming contents, thus later it might not have enough money for a significant content. Here, we study the *self-enhancement* and *self-transcendence* characteristics that fit on the Schwartz value system that an agent can adopt. We employ three categories described in Section 4.1, namely benevolent, pragmatic and greedy. A benevolent agent tries to spend less than its bidding evaluation of a content, while the greedy does the opposite. A pragmatic agent tries to bid as close to its bidding evaluation as possible. Benevolent agents' biddings result in the agents going along with the society choices. Pragmatic agents try to keep a balance between their own privacy decisions and the privacy understanding of other agents entering the auction. Greedy agents always try to dictate their own privacy decisions by bidding higher amounts than they should for auctions.

Again, we first study the effect of money spending according to the defined characteristics on a setting, where a benevolent, pragmatic, and a greedy PANOLA competes with a non-learning agent where all agents consider the content ordinary. Figure 4.7 shows the success rate of each PANOLA against non-learning agents. Since the benevolent agent gives lower bids to let society have their say in the privacy decisions, its desired privacy decisions rarely become the outcome of the auction. The greedy agent starts with winning all auctions, but since it bids higher amounts of units than it gets, its success quickly decreases due to being taxed more, causing the opponent to win almost half of the auctions. Of all three PANOLA types evaluated, pragmatic PANOLA is the most successful in the long run, since it keeps a balanced bidding strategy and learns to bid the necessary amount for winning an auction. This also indicates that humans that are not leaning extremely on either self-transcendence or self-enhancement, and having a more balanced value that puts similar importance to self goals and the society's goals yield better results for collaborative privacy decisions with PANO.

Next, we compare the performance of various spending strategies against each other. For these scenarios, all the opposing agents consider incoming contents as ordinary. The purpose of this approach is to evaluate different spending behaviors in the same conditions. There are three different scenarios where we compare a pragmatic agent against a benevolent agent, a benevolent agent against a greedy agent and a greedy agent against a pragmatic agent over 100 auctions for each scenario.

Pragmatic Agent vs. Benevolent Agent: In this scenario, while pragmatic agent tries to spend according to its actual evaluation for an action, the benevolent agent aims to bid less than its own evaluation. The comparisons according to success and owned budget metrics are shown in Figure 4.8.

As can be expected, the pragmatic agent has much higher success rate than the benevolent agent. The benevolent agent can still learn to bid more in time and win about 20% of the auctions after a while, and it saves a significant amount of budget doing this. This evaluation shows that applying a benevolent strategy makes sense when the outcome of an auction is not much important for an agent, and the preserved budget can be useful in latter auctions that are significant for the agent. The results also show that self-transcendence results in a lower individual success, meaning that




Figure 4.8: (a) Success and (b) Owned Budget of a pragmatic PANOLA against a benevolent PANOLA

people who have this value would give up their privacy for the others.

Greedy Agent vs. Benevolent Agent: In this scenario, agents still have the same importance evaluations, but the greedy agent tries to spend more than its actual bidding aim. Figure 4.9 shows the results of this evaluation.

The results of this scenario have similar, yet slightly closer trends with the first scenario for the success rate. At first thought, even though a greedy agent sounds like a certain success than a pragmatic agent against a benevolent agent, with PANO mechanism the greedy agent depletes its owned budget faster, mostly due to higher taxes




Figure 4.9: (a) Success and (b) Owned Budget of a greedy PANOLA against a benevolent PANOLA

paid. Thus, these two scenarios indicate that against a benevolent opponent, while the bids are greater than the opponent, spending less like a pragmatic agent is a more sensible strategy.

Pragmatic Agent vs. Greedy Agent: The final scenario for money preservation puts the pragmatic agent against the greedy agent. The results of the evaluation are shown in Figure 4.10.

In this scenario, greedy agent depletes its budget quickly, and after the first 30 contents, overall success of the pragmatic agent surpasses the greedy, while still pre-

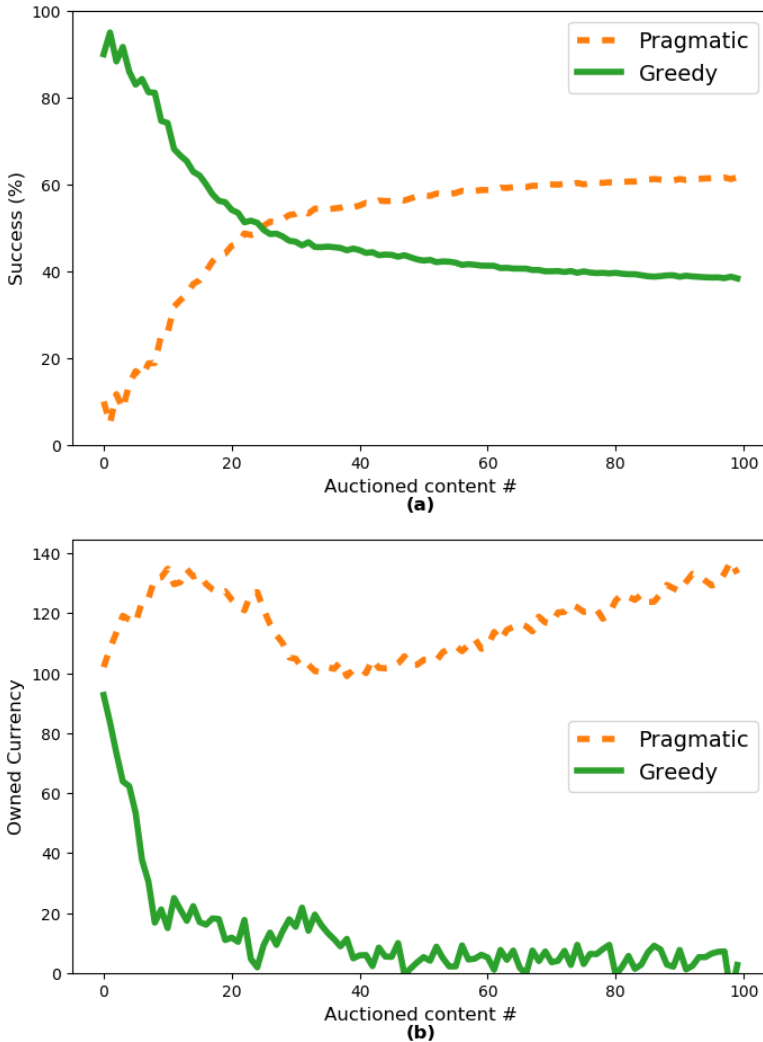


Figure 4.10: (a) Success and (b) Owned Budget of a pragmatic PANOLA against a greedy PANOLA

serving significantly greater amount of budget.

As the result of the three money preservation scenarios, the importance of searching for the smaller amounts to win auctions comes into light. In PANO, bidding more than it should results in getting taxed with great amounts, and in a few auctions to deplete entire owned budget. Even though truthfulness is the best strategy for a single auction as proved in the literature; in continuous auctions, spending less can become a better strategy. However, the bid amount should still be higher than the opposition; therefore, a strategy where the winning bid is slightly decreased for the

future auctions could be resourceful, which would still help an agent to win auctions. Bigger jumps in decreasing bids can be detrimental just because an agent can easily become outbid by the others.

Strategy-wise, we can conclude that greedy characteristic only works for a short amount time and when the agent has enough budget to spend bigger numbers. In these conditions, when a content is significantly private for an agent, the greedy strategy can be employed for a higher chance of win. The benevolent strategy considers the general good of all the participants by letting the other participants have a bigger say by spending less. This strategy enables the agents to preserve money for future auctions which can be for possible significantly private contents, but the desired privacy outcome by the agent itself might not be reached because of the lower bidding for the current content in consideration. The pragmatic strategy can be considered as the most truthful one, since the agent bids from the exact ranges according to its privacy evaluations. Hence, in the longer run, the privacy success of the pragmatic agent turns up better than the other two money spending strategies.

Referring back to human values, the results of our evaluations show that OSN users that are not strictly leaning on self-transcendence or self-enhancement, and rather adopt a balanced valuation where they seek a middle ground between individual goals and goals of others in the society tend to have more success in collaborative privacy decisions using PANO. However, leaning on self-enhancement values might be more beneficial when the piece of content in question is considered very private, and ignoring others' choices. In an opposite case, when the content is insignificant for an OSN user, valuing self-transcendence to go along with the society's privacy choices could result in a higher satisfaction for the whole society.

4.5 Discussion

In this chapter, we have investigated how human values can affect learning strategies for collaborative privacy decisions. We adopt Schwartz values [81] for PANOLA agents, with values that reflect *openness to change*, *conservation*, *self-transcendence* and *self-enhancement*. The results show that human values indeed have effect on how successful an agent can assist an OSN user. For example, a user who values self-enhancement can try to make each collaborative decision to be in line with her individual privacy preferences, while a user that values self-transcendence can give up on her individual privacy choices for the good of the society. In another case, a user can be open to change and prefer experimenting over privacy decisions, while a conservative users would like similar outcomes for privacy decisions over each piece of content. While representing the human values with PANOLA is beneficial to an extent, our results show that sticking with some values do not always yield in better results, either for the individual or the society. Therefore, while assisting for privacy, PANOLA agents should also lead the users they represent to adopt strategies that are not always in line with their values to reach better outcomes for preserving privacy.

The experiments in this chapter indicate that people with specific types of human values can perform better for collaborative privacy decisions, both for the individual success and the privacy preservation for the entire society. For example, having a bal-

anced valuation where a user values self-enhancement and self-transcendence equally results in better outcomes both for the user herself and the society. Moreover, preferring conservation of the current strategies over exploration yields in better results in the long term. We would expect people that share similar values to reach similar outcomes for privacy decisions, since they would act akin to each other. For example, OSN users that value self-enhancement would often reach privacy decisions that are closer to their individual preferences. On the opposite scale, users that value self-transcendence can reach decisions that favor what most of the society wants. This kind of similar behavior indicates that *social norms* could be present in OSNs for privacy decisions, where in a given context, OSN users with similar human values tend to reach similar decisions. Identifying these social norms can enable collaborative privacy decisions to be reached without the need of a complex decision mechanism, where norms can govern decisions to be applied. In the next chapter, we will build upon this idea to investigate if the social norms in an OSN can be identified, and present a mechanism to apply norm-based privacy decisions where agents can simply follow norms instead of employing a complex mechanism to reach collaborative decisions for co-owned content.



5

Privacy Norms in Collaborative Systems

ABSTRACT

Collaborative systems, such as online social networks or Internet of Things, host vast amounts of content that is created and manipulated by multiple users. Co-edited documents or group pictures are prime examples of such *co-owned* content.

Respecting privacy of users in collaborative systems is difficult because the co-owners of the shared content can have conflicting access policies about the content. To address this problem, we have employed PANO for the online social networks domain, which is explained in Chapter 2.

With PANO, when a content is to be shared, all co-owners express their privacy preferences through the mechanism (i.e., by bidding) and the group decision mechanism reaches a decision to enable or deny access to the content.

However, such mechanisms have to be carried out per content, making them impractical for most realistic settings. Even with learning agents like PANOLA, the mechanism would require a learning phase, in which input from users is essential. We argue that rather than employing a group decision mechanism on each content separately, it is more practical to watch for *privacy norms* that emerge in systems and make decisions using these norms, when possible.

This chapter borrows ideas about norm typology from philosophy to represent privacy norms and develops algorithms to compute them in collaborative systems. We propose a mechanism called PRINOR, which identify social privacy norms both for the entire society or the small groups within it, and provides privacy decisions in line with these norms without the need of a complex decision mechanism. The users are represented by software agents, and the agents incorporate an algorithm to either follow the social privacy norms or opt in to make collaborative decisions according to the privacy concerns of the users they represent. We show that when privacy norms are identified correctly, they can enable collaborative systems to respect users' privacy as well as decrease the need to engage in group decision mechanisms considerably.

5.1 Introduction

Many recent software systems are built on the idea of collaborative computing, where multiple users share, manipulate and manage information about themselves as well as others. Online social networks (OSNs) are a prime example, where a user can put up a group picture without explicit consent from individuals in the picture, and others can access to, comment on or even reshare the content, making it more visible to the world. For many users, this means that their personal lives are easily accessible to individuals or companies without them knowing about it.

Even though we are living in a privacy-conscious era with various policies in place to attempt preserving privacy, existing techniques have not been sufficient to detect, let alone handle these privacy violations. The main reason behind this is that privacy has been simplified to an informed consent, where the main assumption is that a user is in control of her data and chooses how to manage her privacy by giving appropriate consent. General Data Privacy Regulation (GDPR) [24] is an important policy, which is based on this idea of informed consent. While GDPR assumes that each user can independently manage the privacy of personal data, the content that exists on collaborative systems, such as co-edited documents or group pictures, do not always belong to a single person. Further, many times content about a user is shared by others, not by the person herself. For example, a co-author of a jointly edited document can send a link of the document to whomever she sees fit. Or a user on an online social network can share a group picture publicly without explicit consent from those in the picture. It is possible that the individuals that are related to the content might have different and possibly contradictory privacy preferences [46, 95]. In these situations, when the sharing party is assumed to own the content, only her privacy preferences will be in effect. However, in many situations, the content might be *co-owned* by others that bring about the content in the first place (e.g., co-editors of a document). Hence, it is not sufficient to allow access to the content by only considering the sharing party's privacy preferences.

GDPR does not address how to tackle the privacy of content that pertains to more than one individual or that is shared by others about the user [12, 74]. We need to think of privacy for co-owned content different than the privacy of personal data since the sharing intentions and privacy preferences of all involved are at stake. In the previous chapters, we have shown that various collaborative privacy management mechanisms exist to tackle this. Negotiation-based agreement techniques, argumentation-based techniques and auction-based techniques are some examples that make use of multi-agent system based approaches in order to resolve collaborative privacy disputes. These approaches are promising when each user is fully aware of her privacy expectations and can actively participate in the decision making whenever an access decision needs to be made. However, this is unrealistic for many systems where huge amounts of co-owned data are shared frequently but many users do not engage in configuring their privacy settings. Even when agents can learn and act on behalf of the users like PANOLA, initial user input is required to capture the privacy requirements of the users they assist. Therefore, PANOLA agents require a learning phase, which was up to 1000 privacy decisions within the OSN, according to our experiments in Chapter 3. Thus, it would be useful to be able to configure

the access settings of a content without explicitly involving all the co-owners of the content into decision making.

Human societies often use norms for decision making [45]. We advocate norms as the basis of collaborative privacy decisions mechanisms. If the systems can identify the existing norms, then decisions can be made using them. This implies that a more complex decision mechanism, such as an auction or a negotiation, would not be required, speeding up the decisions that can be taken considerably. Moreover, norms can also serve another important function. When an individual does not have or cannot formulate her privacy preferences, then the norms of a society can shed a light as to what is appropriate. The user does not have to follow the social norms at all time. If the user does not want to follow the norms or none of the existing norms apply to a given situation, then the system can still use a collaborative privacy management mechanism to make a decision. Contrary to successful access control schemes, such as role-based access control [79] or relation-based access control [38] that mostly apply on content that is owned by a single individual, norm-based privacy decisions enable access decisions on co-owned content, therefore could become an appropriate application for collaborative privacy decisions.

This chapter describes the principles of norm-based privacy decisions and develops an approach named PRINOR where collaborative privacy decisions can be taken based on the *norms* that are generated from the previous privacy decisions in the system. We represent the different privacy norms in OSNs using Tuomela's norm categorization [108] and develop algorithms to identify these norms in a given system. The usage of the algorithms enables users to choose between enforcing personal privacy settings and following the norms in the system. We show over multiagent simulations that when privacy norms emerge, they can be used in place of collaborative decision mechanism that require interactions for each content. Our analysis shows that the variations in the privacy expectations of the users have little effect on the success of PRINOR. We also apply PRINOR on a case study with real-life social network and image content data sets to demonstrate norm emergence and privacy decisions.

5.2 Privacy Norms

We study the representation, emergence and usage of norms in collaborative systems, where a set of users are related to each other through a set of relations types (r_{type}), such as friend, colleague, and so on [38]. Each user can share content that pertains to herself as well as others. A user's privacy preference about a content could depend on the properties of the content as well as the relation types with whom the content is shared. For example, a user might not want her holiday pictures to be shown to colleagues, but might be fine with work pictures to be shown. When the co-owners have conflicting privacy preferences, they need to reach a *privacy decision* that states if and how the content will be shared.

Each user in PRINOR is represented by a software agent, which keeps track of the privacy expectation of its user for sharing content. We represent contents with a content descriptor c_{des} , which is a set of two tuples (x, n) , where x is a context

such as holiday, work, and so on and n is the percentage of how much this content belongs to x . For example, a picture taken at a bar with friends might be represented as: $\{(nightlife, 77\%), (leisure, 12\%)\}$. The context information might be available in the system but it could also be derived automatically as we will explain in Section 5.6 through a software that produces tags and confidence intervals. Depending on the content, the set might have more tuples. We do not require the sum of the percentages to be equal to 100% since the content may be highly relevant to multiple contexts making their sum way over 100%. Alternatively, we may not have enough evidence to associate a content with contexts; hence the sum may be less than 100%. For each content c , we also specify the set of co-owners c_{own} , whose privacy is possibly being affected by the content and thus should have a say about content's privacy decision. In general, if the content is a picture, c_{own} could consist of users tagged in the picture or if the content is a co-edited document, it could be the co-authors. We assume that this set can be retrieved from the system as is the case with most collaborative systems.

PRINOR contains norms to capture the privacy preferences. Informally, privacy norms capture the common behavior for accessing a particular type of content with a particular set of users. In most domains, it is generally assumed that when the actions of the agents are in conflict, norms that are fully applicable to all the agents can be found. A prime example is the well-studied traffic domain [65], where norms such as driving on the right side of the road might emerge because it can be observed that mixed usage of sides leads to accidents. However, norms have to be rethought in the case of privacy. Privacy norms, by definition, are different from other norms because it is extremely difficult to find privacy norms that could satisfy the expectations of all of the agents. That is, there is no one right norm to make everyone happy. Whereas in a domain such as traffic, an agent benefits from obeying an established norm; in a privacy related domain, complying with a social norm might harm the privacy of agents, depending on their privacy requirements. Therefore, norms should be allowed to emerge at different levels (e.g., norms in a group of users, norms for a specific relation type, and so on) and the norms should be evaluated continuously to find out if they still fit to the expectations of the population. Moreover, an agent should always be allowed not to follow a norm so that privacy preferences of individuals and minority groups can still be respected.

Tuomela [105] categorizes norms as social and personal norms. Social norms are formed according to the behavior of the society, and can be sanctioned if one does not comply with them. Personal norms are related to individuals' comprehension of the environment, and their beliefs about which actions are right or wrong within the society. Tuomela further divides social norms to r-norms (rule norms) and s-norms (social norms), and personal norms to m-norms (moral norms) and p-norms (prudential norms). We adopt this classification to model privacy expectations as norms and formally represent it similar to existing formalisms [9, 65, 108], such that a set of preconditions determine the activation of a sharing action to be taken. We also aim to handle context-based privacy preferences [10] with our norm definitions. Since our focus is more on the emergence of norms rather than their violation, we do not include norm sanctions explicitly [80]. We consider social norms to be governed by an overseer mechanism (e.g., an OSN provider), while personal norms are handled

in a distributed manner.

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

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

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

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

5.3 PRINOR

The norms associated with the norm types described in Section 5.2 are all stored in respective norm bases. Initially, each agent has a personal *m-norm* base, which can only be updated by the agent itself. An *m-norm* base can be thought as the privacy policy of the agent. For now, we assume that the *m-norm* base for an agent does not change over time. At the beginning, the collaborative system itself has a single *r-norm* base that contains all the laws of the system. The norms in the *r-norm* base are stored and updated by the system provider itself. Again, we assume that the *r-norms* are static and do not change over time.

S-norm base contains the social norms in the system. These social norms emerge

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

PRINOR works as follows: When an agent wants to share a content, which is co-owned by other agents, the uploader agent checks if it is desirable for all the co-owners to share the content, considering the norms. This is done by considering the type of the content and the relationship with other co-owners. Since four types of norms are in effect, there can easily be conflicts among various norm-bases. For example, an agent's *m-norm* might permit sharing a content publicly, whereas the *s-norm* in the system might prescribe otherwise. This calls for an ordering of norm bases. Dechesne *et al.* [30] show that there are several individual characteristics that affect the decision process of the individuals, such as compliance with the law, abiding to social conventions or behaving according to individual preferences. An individualistic agent might first check its *m-norm* base and refer to other bases only if there are no related norms in this base. A social agent can prefer to put *s-norms* in front of *m-norms* while a law abiding agent always places *r-norms* at the top. An interesting choice question comes up with *p-norms* and *s-norms*, since they are both in the social context, while the former only includes a specific set of agents that the agent directly has a relationship with. In this work, we assume that *p-norms* always dominate *s-norms*, since norms within direct relationships represent more precise behavior than the norms emerging from a community which is formed by indirect relationships (i.e., agents that do not have a relationship, but are present in the same OSN community) and that *s-norms* dominate the *m-norms* since we are interested in understanding the benefits of making privacy decisions using societal norms. Using this ordering, the uploader agent checks its *r-norm*, *p-norm* and *s-norm* bases to see if a norm matching with the content descriptor exists. If so, it is applied. It might be the case that none of the norms in the norm bases are applicable. If so, the agents engage in a decision mechanism, such as auctions or negotiation, and the final decision is made according to the chosen collaborative privacy mechanism. When agents engage in a decision mechanism, they use their *m-norm* bases to reveal their valuations. If such a mechanism is used, then the outcome of the mechanism also updates the *p-norm* base of the co-owner agents and *s-norm* base of the OSN, where new possible norms can be formed for future incoming co-owned content.

Figure 5.1 depicts how PRINOR works when a decision is to be made for an incoming content for three agents, Alice, Bob and Carol, who have a friendship rela-

<p>m-norm Base {Alice}: m<friends,{<work, 100%>},deny> m<colleagues,{<beach, 100%>},deny></p> <p>m-norm Base {Bob}: m<friends,{<work, 100%>},deny> m<family,{<nightlife, 100%>},deny></p> <p>m-norm Base {Carol}: m<friends,{<leisure, 80%>},access></p>	<p>p-norm Base {Alice, Bob & Carol}: p<{Alice,Bob},{<work, 100%>},deny></p> <p>s-norm Base s<family,{<work, 86%>},deny> s<friends,{<leisure, 82%>, <scenery, 52%>},access></p>	<p>Content #X content descriptors: {<work, 86%>, <leisure, 4%>}</p> <p>Content #Y content descriptors: {<leisure, 72%>, <scenery, 21%>}</p> <p>Content #Z content descriptors: {<nightlife, 77%>, <party, 12%>}</p>
---	--	--

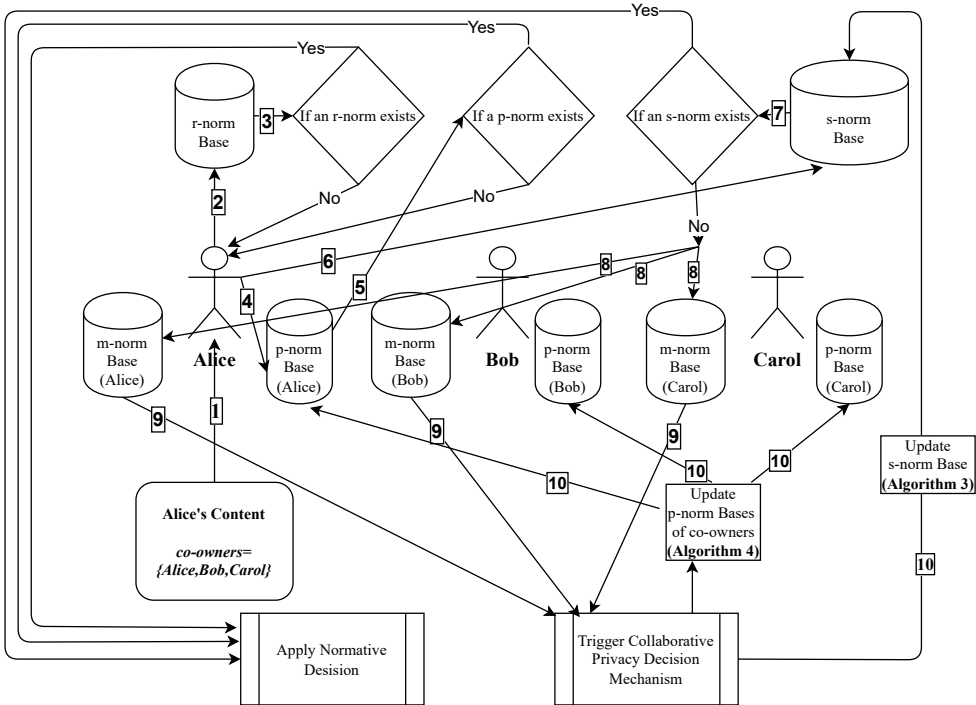


Figure 5.1: Normative decision mechanism process for an incoming content co-owned by three agents.

tion. The legend on the left side describes the norms and the contents relevant in the system. The numbers of arrows indicate the order of actions. Alice wants to share the contents but the contents are co-owned by all three agents. We give various examples of how PRINOR would work on an incoming content. In our examples, we exclude r-norm base checking for brevity, since it is only applied simply when a content type is forbidden by the OSN provider and always checked initially by all agents to not receive possible sanctions.

Example 9 The incoming content *X*, which is mostly related to work context, according to the content descriptors. As the uploader, Alice checks her *p-norm* base, where all previous normative privacy mechanism decisions by every subset of these

three agents are stored. *p-norm* base includes a fitting *p-norm* established between Alice and Bob, with *deny* action. Since this norm would be in effect in the greater group as well, the agents do not share the incoming content.

Example 10 The incoming content is Y, and mostly related to leisure context. Since Alice does not have a related *p-norm* in her *p-norm* base, she checks the *s-norm* base and finds a similar *s-norm*, where the content descriptor indicates 82% relatedness with leisure context. Alice, Bob and Carol can comply with this *s-norm* to share the content according to *s-norm* base, since they are indeed in friendship relation.

Example 11 For incoming content Z, Alice does not have an established *p-norm* in her *p-norm* base. Content Z does not fit into *s-norms* in the *s-norm* base, either. Therefore, collaborative privacy decision mechanism should be triggered, and the decision should be made according to agents' *m-norms*. Since this is a mechanism based decision, *p-norm* base of Alice, Bob and Carol is updated with the current decision. The *s-norm* base of the OSN registers this decision to be used for the generation of norms.

5.4 Generating Norms

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

5.4.1 Identifying S-norms

Recall that each co-owned content in the system requires a privacy decision according to their contextual properties, and the outcome is to enable or deny access to the content with a set of relationship types. Given a set of such decisions, Algorithm 3 clusters the decisions to identify potential *s-norms*. Essentially, the algorithm places all the content over a multidimensional space according to their descriptor and the relationship type of the co-owner agents. This space contains all the decisions considering its various properties as dimensions. Then, we cluster this space such that each cluster contains content that have similar attributes. Finally, the clusters can be checked for being a possible *s-norm*. Since the evolution of social norms depend on many factors, continuous update of *s-norms* is essential to capture the current state of social normative behavior in the environment [80]. Therefore, the algorithm is run periodically in order to find out about new emerging norms or exclude norms that became obsolete over time.

OSNs enable users to continuously share tremendous amount of content. In a real life application, clustering every content in short periods would be infeasible, since it would require massive computing power. Therefore, a simple clustering algorithm which is sufficient enough to distinguish between contextual properties of content, in our case, the dimensions of the descriptors, is essential. In light of this intuition, we employ *k-means* algorithm to cluster content and then check all the clusters for normative behavior. *k-means* is a clustering method where n number of elements in a unidimensional or multidimensional space are partitioned into k clusters, where each element is assigned to the nearest mean of the elements in a cluster [99]. Determining the number of clusters is difficult. Rather than having a fixed number, we adjust it as needed. More precisely, in Algorithm 3, we start with a small number of clusters containing large amounts of content, hence we define a small k value resulting in a big n value. As a heuristic to determine normative behavior within a cluster, we use qualified majority to check the privacy decisions for the content within. According to qualified majority heuristic, we consider a cluster a candidate for being normative, if at least 66% of the privacy decisions are the same for the content in the cluster. The threshold amount 66% is picked for qualified majority, since it is commonly adopted in many domains such as law and voting mechanisms. If a candidate normative cluster is found with the initial k , it is saved to the *s-norm* base and the content within is removed from further calculations. For the remaining clusters that does not show normative behavior, k is increased and new clusters are formed to check if normative behavior emerges with smaller number of content within clusters. This approach continues until a threshold for the minimum number of agents in a cluster is reached, and the algorithm stops at that point to save the rest of the clusters as *non-normative*.

When applying naive *k-means* clustering, each content can be placed into the closest cluster, because all content is assumed to have the same dimensions. However, shared content in real life would have differing contextual properties. Since we take each contextual property as an additional dimension, the number of dimensions might become high. Moreover, many content would not have a common contextual property. A simple approach would be to still consider all possible contextual properties as separate dimensions, and assign the value of zero if a content is not related to this content descriptor. This would make the space rather sparse. With a large number of dimensions, this can easily become infeasible since each content would have many dimensions valued at zero. As a result, the clustering can yield very distant clusters each containing only a small number of content. To resolve this, we propose a dimension reduction for the domains with a large variety of content descriptors. With this reduction, for a privacy decision of an incoming content, only the previous content that share descriptors are taken into consideration for clustering. Thus, the only dimensions required for clustering would become the incoming content's descriptors, which would significantly reduce the computation required for finding *s-norms*. For example, if there are 1000 previous content decisions are present in the *s-norm* base, while only 50 of the decisions share descriptors with the incoming content, we would not consider the remaining 950 content that do not share any descriptor, since the new content does not share any contextual property with them. In this case, the computation would only require the use of 50 previous decisions that share descriptors.

Another aspect to consider for social norms is the changes in the behavior of the

Algorithm 3: Generation of s -norms

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

```

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

society over time. As the time passes, the privacy preferences of the people change, which also might cause some norms to become obsolete while new ones emerge. Thus, we employ an aging curve [117] for privacy actions, denoted as $R = e^{-t/S}$. Here, R is the *retrievability* of the privacy action, while t is time passed since the decision was taken and S is the *stability* of memory. According to this equation, a recent privacy decision would take a value closer to 1, while over time it's value would be close to 0, and the speed of aging is dependent on the S value. Let us consider two examples of privacy decisions, one taken an hour ago for sharing a content while one taken 1000 hours ago for not sharing. If we define the S value according to a calendar month, hence, 720 hours, the first one would have a value of ~ 1 for retrievability, while the second is ~ 0.25 . If these were the only privacy decisions in the system, the naive qualified majority calculation for s -norms without considering the aging curve would have given 50% for sharing and 50% for not sharing a similar content. With aging curve in place we give the calculated R values as weights to the privacy decisions. Hence, the two example privacy decisions would result in 80% for sharing and 20% for not sharing, since the recent privacy decision is considered more significant for capturing current social behavior.

For each periodic call of s -norm base update, the minimum cluster count parameter (mk), the minimum size threshold parameter for a single cluster (t), all the previous privacy mechanism based decisions ($pDec$) and stability value (S) are taken as input for Algorithm 3. The algorithm starts with calculating retrievability values of all previous privacy decisions in s -norm base according to the stability parameter (lines 1 and 2). Then, for each item in $pDec$, a temporary list of clusters are assigned

with k-means algorithm, where all items in $pDec$ are clustered into mk clusters. In line 5, a for loop begins, which checks the temporary cluster assignments, and determines if the cluster shows a normative behavior (i.e, qualified majority of the privacy decisions are the same), or the size of the cluster is below t value. If one of these conditions is satisfied for a temporary cluster, the cluster is added to $cList$ in line 9 and all the items of the cluster are removed from $pDec$, ending the iteration. If there are still remaining items in $pDec$, another iteration starts to determine new clusters, until all items from the initial $pDec$ are assigned to a cluster in $cList$ output.

5.4.2 Deciding to Follow an S-Norm

After s -norms are identified with Algorithm 3, they are stored in the s -norm base. Whenever an agent is making a privacy decision, it will check the s -norm base to see if any of the s -norms are applicable. If so, then the agent needs to decide if it would want to follow it.

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

The *majority percentage* for the s -norm privacy action is provided by Algorithm 3, which requires no further computation. To compute *contextual similarity*, we place the incoming content in the cluster and compute the Euclidean distance of every content descriptor dimension to the center. Then we do the same for the content in the same cluster that is furthest from the center. With the second distance, we normalize the first distance in a way that the furthest content would give *contextual similarity* value as 0 and the center itself would be 1. For example, in a single dimensional context, if the distance of the furthest content to the center is computed as 4, and the distance of the incoming content to the center is computed as 1, *contextual similarity* would be $((4 - 0) - 1)/4 = 0.75$, which means the content is strongly related to the cluster contextually. To compute the *decision similarity* metric, agents check if they have any m -norms for a similar type of content. If they do, then the *decision similarity* is simply the number of these moral norms divided to the number of all m -norms stored by the agent. This affects the decision for considering the incoming content

normative positively, if the privacy decision of the considered *m-norms* are the same with the *s-norm*'s majority decision. If not, the effect of the *decision similarity* becomes negative.

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

$$SD = \begin{cases} \alpha * (CS * MP) + (1 - \alpha) * DS, & \text{if } > 0. \\ 0, & \text{otherwise.} \end{cases} \quad (5.1)$$

Equation 5.1 divides the defined metrics into two parts according their relation with each other. Contextual similarity of the incoming content for the cluster, and the majority privacy decision of the cluster are closely related to each other, hence we multiply them with each other, and also multiply this with α for weighting these parameters' importance. If the contextual similarity is high and it is a strong norm with a high majority percentage, this would result the multiplication to have a value that is closer to 1, which yields in an *SD* value that strengthens the suggestion of using the *s-norm*. Decision similarity is related to differences of individual choices against the society, therefore it is weighted with $1 - \alpha$. With an α value closer to 1, the *SD* is more reliant on the society choices, while a value closer to 0 would give more importance to the similarity of individual privacy requirements with social norms. Referring back to the examples in Figure 5.1 of Section 5.2, we give an example below to depict how agents can use the *SD* metric to decide complying or denying social norms.

Example 12 Consider Content Y from Figure 5.1 as the incoming content uploaded by Alice, and *MP* of the related cluster is 84% for not sharing action, given by Algorithm 3. For the calculation of *CS*, the content descriptor has two dimensions. Let us assume the furthest content of the same *s-norm* cluster has a content descriptor as {<leisure, 32%>, <scenery, 13%>} and each context type has the same importance. For the leisure contextual dimension, the furthest content of the cluster has 50% (82%-32%) and 39% (52%-13%) distance for the scenery context. The same values for Content Y are 10% (82%-72%) and 31% (52%-21%), respectively. Therefore, *CS* value is the mean of (50%-10%)/50% and (39%-31%)/39%, equalling to ~50%. Alice does not have any *m-norms* related to the *s-norm* in consideration in her *m-norm* base, which is consisting of two *m-norms*. Thus, the *DS* value will be 0 (0/2). If the α is given as 0.8, *SD*, which is the likelihood of complying with the norm would be calculated as $(0.8 * (0.84 * 0.5) + 0.2 * 0)$, which would prescribe Alice to follow the norms with a probability of ~33%.

5.4.3 Identifying p-Norms

Prudential norms are the second type of societal norms in our mechanism. The *p-norms* only bind the users in the group and not the society as a whole. Hence, we require that *p-norms* are kept and updated separately by each agent (rather than by the OSN provider as was the case with *s-norms*). Essentially, *p-norms* represent previous collaborative privacy mechanism based decisions of co-owner agents for a content. To keep privacy requirements simple, agents only classify *p-norms* according

to the major content types (i.e content type with the highest relatedness value). In addition, agents keep track of the co-owner IDs, since *p-norms* are the norms that emerge between specific sets of co-owners. Algorithm 4 shows how an agent generates a *p-norm* after deciding on an incoming content.

Algorithm 4: Generation of *p-norms*

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

```

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

```

After an incoming content where a privacy decision is required, Algorithm 4 is triggered by each co-owner agent of the content to generate a *p-norm*. The inputs include the major content type (*mct*) of the content, which defines the highest valued content type. *aDec* contains previous privacy decisions of co-owners, including the decisions made by subsets of the co-owners. This enables the algorithm to propagate previous privacy decisions of smaller subsets of co-owners to the entire set of co-owners. Since a subset of co-owners might not fully represent the behavior of a bigger co-owner group, we introduce a difference parameter (*d*) in the algorithm, which enables the system to adjust the impact of previous privacy decisions with different size of subsets of co-owners. The algorithm starts with assigning counts of each action type possible for a privacy decision as zero (line 1). Then for each item in the *p-norm* base of an agent, the algorithm counts the previous privacy actions, where the major content type is the same as the current content, and all the co-owners of the item in *p-norm* base are elements of the co-owners set of the content. In line 4, the difference between the size of co-owners of *p-norm* base item and the size of co-owners for the content in consideration is computed. For example, if the content

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

5.5 Evaluation

In this chapter, we have established a norm-based privacy decision mechanism that agents can follow to reach collaborative decisions. The system identifies social norms that are present in the OSN (s -norms), while the agents decide on if to follow these norms or decide by themselves. Moreover, we acknowledge that societal norms in some groups within the society can also emerge, which we identify as p -norms. With an ordering of various norm-types, agents can make privacy decisions either with the help of the norms, or still decide with a privacy mechanism that takes each co-owner's privacy requirements into account. In order to evaluate the performance of norm-based privacy decisions, we answer the following research questions:

- **RQ. #1:** Do s-norms and p-norms emerge over time and if so, what percentage of privacy decisions are taken by these norms?
- **RQ. #2:** Do the norms that emerge enable agents to make correct privacy decisions?

We evaluate PRINOR in a multiagent simulation environment that we developed in Java¹. Each agent in the simulation represents a user. The users, and thus the agents are related to each other through one relationship. Each agent is uniquely identified with an identifier. Each agent has a set of m -norms that are generated automatically. Each content in the OSN is assigned a descriptor. In real life, this information would come from the features or tags of the content. Here, we assume that the descriptor is available. For n number of content type categories, a content

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

is placed in an n -dimensional space which enables the mechanism to both find out similar content types and match privacy requirements of agents with the content in consideration. In addition, a content has a set of co-owners, which are the agents that are within the OSN that have some of their private information represented in the content.

We include 100 agents and 10000 contents for each of our simulation runs, where each content is randomly assigned to 2 to 5 co-owners, and a descriptor with 4 elements, while each element is a two-tuple with a context and a value between 0 and 100, representing the significance of the content to the given type, 100 being the most. We represent each agent's privacy requirements with m -norms, while the simulation checks the evolution of p -norms and s -norms. We exclude r -norms from the simulations as our focus is on the correct emergence of p -norms and s -norms. On each simulation, one content is introduced to the mechanism sequentially. First, the societal norms are checked to reach a decision. If relevant societal norms are not present, then the decision is made according to the m -norms of the agents. For m -norm based decisions, our current mechanism allows us to employ different mechanisms such as auctions, negotiation or argumentation. However, these mechanisms require rather complex computation. In order to keep computational complexity low, we employ majority voting as the collaborative privacy mechanism in our evaluations. With this simulation setup, we aim to answer our research questions about norm emergence (RQ. #1) and the correctness of norm-based decisions (RQ. #2).

5.5.1 Emergence of Social Norms

In a pioneering work on norm emergence, Sen and Airiau [82] show that norms emerge even when the population size and heterogeneity vary. Following this, we introduce a *homogeneity* variable to capture how much of a society has similar privacy understandings. In our approach, if the homogeneity of the society is 0, then all the agents in the population can have different privacy choices. We run our simulations for investigating emergent social norms with different levels of homogeneity. We achieve this by making a subset of agents having the same action type for a given type of content in their m -norm base, while the rest is assigned a random type of action for their m -norms. Our homogeneity levels are 0%, 10%, 20%, 30%, 40%, 50%, 75% and 100% respectively, 0% representing full random behavior and 100% full consensus. The reason of having bigger margin between the latter three levels is that societal decisions are almost similar when homogeneity percentage is bigger than 50% in the network. The simulation starts forming s -norms using Algorithm 3 after 1000th content shared in the OSN and reruns it after every 250 content for updating the s -norm base of the OSN. To compare how agents decide to follow the norms, we evaluate four different setups. In the first setup, agents follow s -norms at least when qualified majority percentage for a single privacy action is satisfied. The other three setups employ the SD formula, which gives a likelihood value of following the norms for agents, with three different α values, 0.8, 0.5 and 0.2 respectively. For each homogeneity level combined with each of the four setups, we run 5 simulations and measure the percentage of decisions taken with m -norms, p -norms and s -norms.

Figure 5.2 plots the percentage of decisions that are taken by m -norms, p -norms

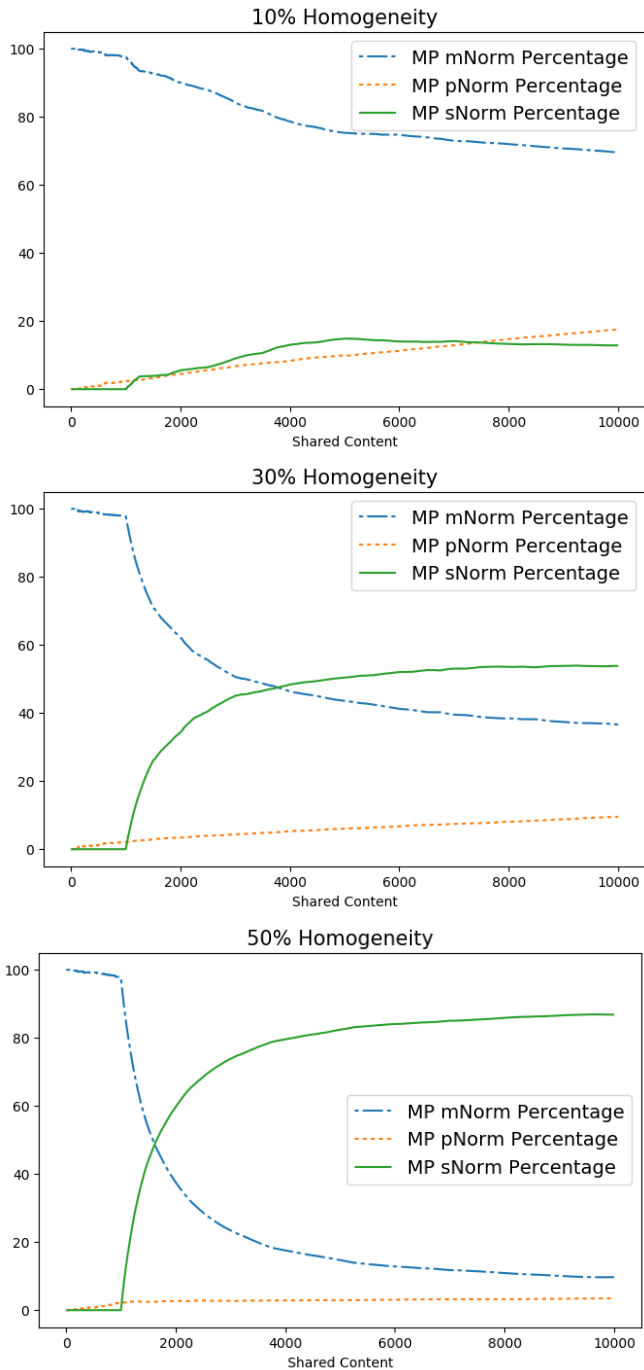


Figure 5.2: Percentage of norm types over different levels of homogeneity for qualified majority s-norm decision.

and *s-norms* as new content is introduced to the system for populations with 10%, 30% and 50% homogeneity and when the *s-norm* decisions are made according to qualified majority percentages of *s-norm* clusters. For 10% homogeneity, 12.88% of all decisions were made with *s-norms* while 17.54% of all 10000 content is decided according to *p-norms*, without the need of triggering the collaborative privacy decision mechanism with *m-norms*. This can be seen as a significant improvement, since our norm based method reduces the need to trigger a decision mechanism by ~30 percent, even when a tiny fraction of the society behaves homogeneously and the amount of co-owned content is sparse. The sparsity comes from having 100 agents randomly assigned as co-owners of 10000 content, since the same subset of agents can only have a very limited number of content with the same major relation type and content type. Therefore, building up *p-norm* base of every agent becomes a difficult task with the limited previous knowledge about the co-owned content with the same subset of related agents. With 30% homogeneity, more than 63% of the decisions can be made with *p-norms* and *s-norms* and with 50% homogeneity, the necessity of applying a privacy decision algorithm with *m-norms* goes below 10%. Our results show that even if a small amount of agents in a system act similar instead of randomly behaving, social norms can emerge and effectively be used for collaborative privacy decisions.

Even though deciding only according to qualified majority decisions for *s-norms* reduce the need of a collaborative decision mechanism significantly, the emergent norms might differ from the privacy understanding of individual agents. Some agents might act differently than the society, therefore applying social norms might create privacy decisions that the agents would not want to achieve by themselves. To account for this, we introduced a likelihood to follow *s-norms* formula (*SD*) in Section 5.4. We evaluate the *SD* formula with three α values and plot the results for 30% homogeneity in Figure 5.3. We omit the rest of the evaluations with different homogeneity for brevity, since all levels show similar behavior in comparison with qualified majority based decisions and this homogeneity reflects the real life social behavior more than the both sides of extreme levels.

Recall that when α is high, agents assign a high weight to a given *s-norm* but value their own *m-norm* about a content less, if such a norm exists. Accordingly, one would expect that with high values of α , more decisions would be taken with *s-norms* and with low values, the number of decisions would decrease. Our results confirm this. The results show that when α is set to 0.8, the decisions based on *s-norms* takes up two thirds of all, and with *p-norms* the total norm based decisions constitute ~75% of the privacy decisions. The number of *s-norm* based decisions decrease with $\alpha = 0.5$ setup almost to the number in qualified majority percentage setup. The number of decisions is even fewer when α is assigned as 0.2, but still reduces the need of mechanism based decisions to less than 46%, when combined with *p-norms*. However, this decrease can still be beneficial for the agents, since they ensure that the applied social norms are in line with their own privacy understanding, while rejected norms are quite different than theirs, since the lower α value enables only accepting the norms that are contextually very similar to the agent's own privacy requirements. This brings up the *correctness* of the applied social norms into question, which we will investigate in the next subsection.

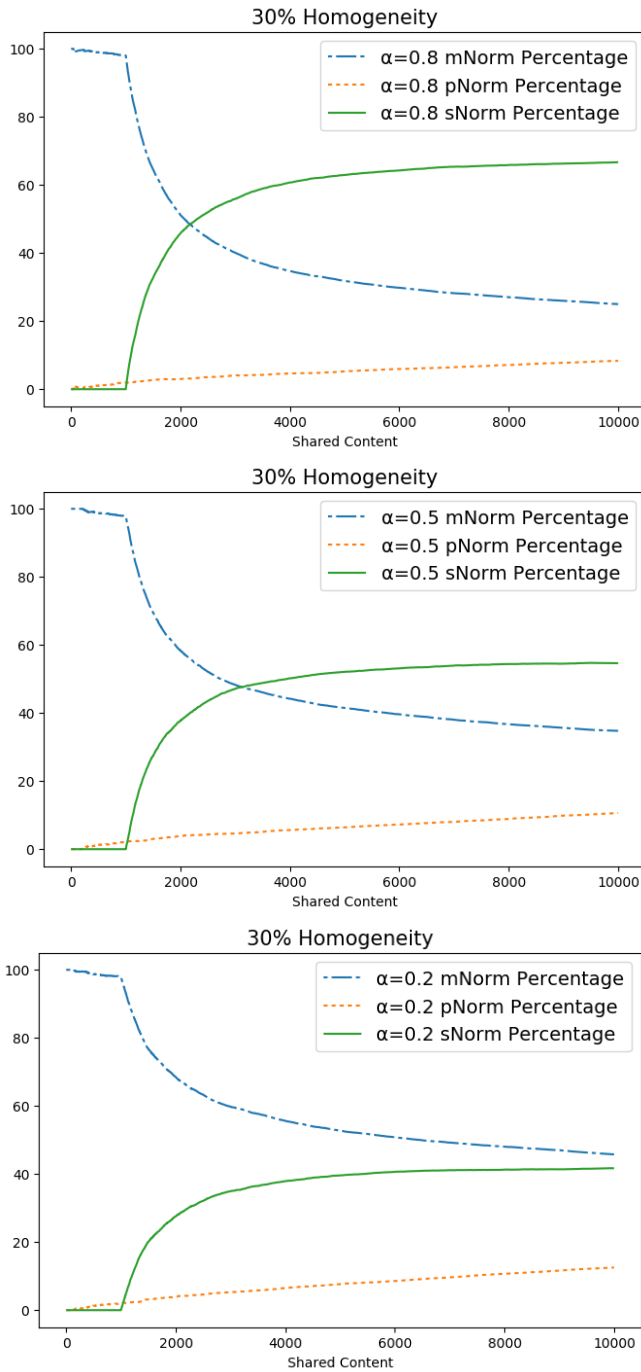


Figure 5.3: Percentage of norm types over 30% homogeneity based on likelihood of following the norms function for various α values.

5.5.2 Correctness of Social Norms

Usage of norms decrease the need of a complex privacy decision mechanism, but do they lead to correct privacy decisions? We measure correctness by comparing norm-based decisions with collaborative privacy decision mechanism results. If the outcome of the norm based decision is the same with what the mechanism would give, we consider it as a correct decision. Since our current setup enables the simulation to evaluate both emergence and correctness of the norms within the same run, we investigate the correctness aspect of PRINOR with our multiple runs for various homogeneity and α values executed for Section 5.5.1 and present our findings about it in this subsection.

Table 5.1 shows the percentage of *s-norm* and *p-norm* decisions over all our setups with various homogeneity levels and agent decision types to follow *s-norms*, along with their correctness ratios. An immediate result is that in any setup, decisions made using *s-norms* are at least 75% correct (HL=0%, $\alpha=0.8$). The percentage of correct *s-norm* assignments increases with higher homogeneity levels, and end up at 100% when all agents in the community are homogeneous in their privacy actions. When we compare different *SD* setups, we observe that the highest correctness percentage comes with $\alpha = 0.2$ parameter. This is an expected outcome since with lower α values, the agents mostly follow the social norms when they are in line with their own privacy policies. $\alpha = 0.8$ setup with the *SD* metric performs the best with lower homogeneity levels to reach a high number of emergent norms, while keeping a reasonably high correctness ratio. Qualified majority setup has the highest *s-norm* percentages with the highest homogeneity levels, since almost all the agents behave the same.

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

Considering the emergent *p-norms* according to Table 5.1, it is seen that *p-norm* assignments are almost always correct. They only depend on subsets of the decisions made by the same co-owners for the similar content types, and unless these agents change their behavior within their own co-owner groups, the *p-norm* based decisions would be the same as the decisions taken with the privacy decision mechanism. However, *p-norm* based decisions are usually a small part of all decisions, mostly due to sparsity of contents over different co-owner agent groups. *p-norms* require more mechanism based privacy decisions to emerge, and with emergent *s-norms*, *p-norms*

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

Table 5.1: Correctness percentages for various levels of homogeneity (*HL*) and *s-norm* decision (*SD*) types.

build up slower than *s-norms*. Note that identifying *s-norms* requires a centralized location that holds the privacy decisions of the society. The OSN itself could provide this location and identification service. If there is no centralized location to enable identification of *s-norms*, PRINOR can still work with *p-norms*, as these are identified in a distributed manner and capture the norms of smaller groups of agents.

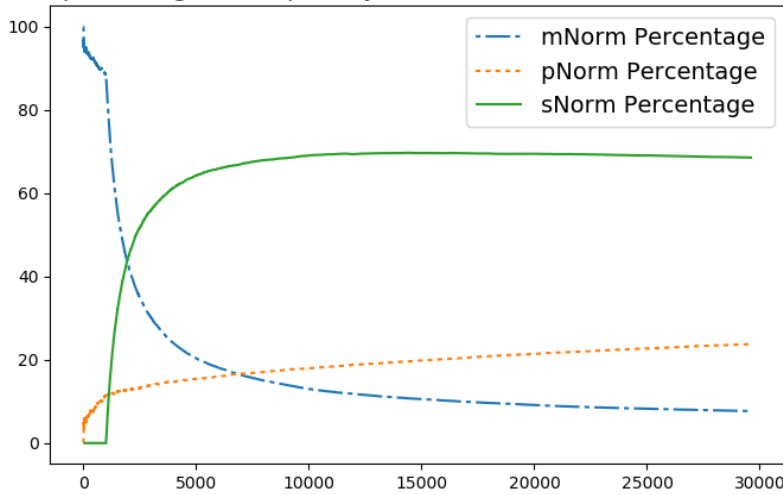
5.6 Case Study for Real-Life OSNs

Our simulation results show that with a given set of contextual features, privacy decisions for OSN content can be made successfully, both with *p-norms* and *s-norms*. However, real life OSNs pose further challenges. First, identifying contextual features of shared content is usually difficult, since users who share the content do not provide these properties. Thus, the contextual identification phase should rely on either the OSN provider or software agents that represent users. Second, some users would potentially have closer relationships and co-own more content than users with limited connections, necessitating the social network to reflect this. We tackle these with a case study by making use of two real-life data sets, namely *SNAP* [59] and *PicAlert* [119] and demonstrate the applicability of our approach.

PicAlert data set consists of images that have been annotated as private or public by study participants while *SNAP* data set contains friendship networks of Facebook users, including their bidirectional relationships, their circles, and anonymized personal features. Since our focus is on norm-based content privacy decisions by software agents on behalf of OSN users, our setup in this case study employs *PicAlert* for defining the contents and their privacy labels; and *SNAP* for defining agents and their network. We extract the contextual properties of *PicAlert* data set with an automated feature extraction tool, named *Clarifai* [118]. We assign four automatically generated tags to each *PicAlert* content as content descriptors and assign *SNAP* agents as co-owners.

Our setup in this case study is as follows: First, we generate all the possible *circles* between the agents in *SNAP* network. We define a circle as a relationship bond between multiple agents, where every two agents in the circle have an established relationship with each other. With this definition, each circle containing more than two agents would have subset circles, since all the subsets of a circle would still be a circle. Our second step is to pick content from *PicAlert* data set and to allocate the content to a circle. We run *Clarifai* API to get contextual tags for the image and assign the four most related tags of the image as content descriptors. Co-owners of a content are picked randomly from all possible circles. Second, we pick a number of content shared in the network and use them to generate m-norms for the co-owners of these shared content. We do this by considering the unique human decisions made for the selected content (available through the data set) and matching the humans with the content co-owners. For example, if a human mostly has share decisions for a contextual tag, an m-norm with share action for the given tag is created for the matched co-owner. As a result of this, some agents acquire m-norms for possible future decisions, while the remaining agents do not establish a privacy understanding. This second category need to rely on emerging p-norms and s-norms to make decisions

Norm percentages over privacy decisions (trained with 1000 content)



Norm percentages over privacy decisions (trained with 5000 content)




Figure 5.4: Percentage of norms over privacy decisions.

to protect their privacy.

The particular *SNAP* data set that we use has been extracted from Facebook and contains 347 users. These users have 975347 possible circles with two or more people. With the *PicAlert* data set, we generate content descriptors for 29864 content. We experiment with two setups of different initial content size. In the first setup, we extract 1000 content and generate m-norms to the co-owners of them. In the second setup, we increase the set to 5000 content. We repeat each setup five times, where different parts of the data set are picked as initial content. We display our results for the percentages of all norm types over incoming content in Figure 5.4 and correctness

ratios after each 5000 privacy decisions in Figure 5.5.




Figure 5.5: Correctness of p-norms and s-norms after each 5000 privacy decisions.

Figure 5.4 shows the percentage of norms for privacy decisions that has been taken in our system, from the first content considered to the last one. For both scenarios with different numbers of initial content, s-norms quickly start to emerge after the first 1000 decisions are made with either m-norms or p-norms. Around 5000 access decisions, s-norms converge to an amount of 70% with 1000 content used in training

and 75% with 5000. Since these 5000 access decisions are for the content co-owned by one of all possible 975347 circles, the s-norms provide a majority of privacy decisions without the need of another mechanism. According to Figure 5.4, we can conclude that even when most of the agents do not have any established m-norms, our approach is still able to form s-norms and p-norms to make privacy decisions. With a larger initial content set, s-norm ratio improves, because agents would usually access s-norms earlier, and would not have the need to make collaborative privacy decisions that eventually result in forming p-norms. Figure 5.5 plots the percentage of correct and incorrect p-norm and s-norm decisions after every 5000 privacy decisions. When m-norms are formed from 5000 content, incorrect s-norms are only 1% of the decisions while p-norms are almost always correct. With 1000 content training scenario for m-norm formation, incorrect decision percentage gets slightly higher for both p-norms and s-norms, with s-norms reaching 3.5%, which can still be considered as a small part of the entire decisions. Thus, both setups achieve correct privacy decisions most of the time.

In order to inspect whether the found social norms are in line with actual OSN users, we analyse some examples of the emerging norms to grasp the intuitions behind them. Combination of the following tags *baby*, *child*, *cute* and *little* generate social norms to deny access to users, as they are commonly seen as private. *Adult* tag, combined with tags such as *man*, *woman* or *person* again generate norms that deny access to other users. On the other hand, tags like *city*, *no person* and *nature* prescribe social norms that deem the content containing them as public, enabling access to other users. These examples indicate that the found social norms resonate with privacy understanding exhibited by humans. When a piece of content does not indicate any information about a human, the chances of a privacy violation is intuitively lower than a piece of content that show a baby, or grown people in some settings such as a bar or a beach. Therefore, we would expect the emerging norms to follow similar cases, which yields results as expected in our case study for the investigated contextual tags.

5.7 Discussion

In this chapter, we have shown that identifying social norms that emerge in OSNs can be used for collaborative privacy decisions with our mechanism called PRINOR. Even without the use of a complex decision mechanism such as negotiation, argumentation or auction-based mechanisms, the privacy outcome for co-owned content is successful in preserving privacy. Even when the society acts heterogeneously, PRINOR can still provide norm-based privacy decisions that fit well with most of the co-owners' privacy requirements.

5.7.1 Overview

Engineering privacy respecting methods for ubiquitous information systems has become crucial as the amount of online information is huge [44, 57, 87]. An important line of research focus on the specification and compliance of individual privacy preferences. Barth *et al.* [13] present a logic framework, formalizing aspects of contextual

integrity and compliance with privacy norms. Barth *et al.* [14] study privacy for business processes, investigating if workflows would lead to data exposure or can verify that the privacy goals are achieved. Basin *et al.* [16] develop a monitoring tool to check policy compliance by employing first-order temporal logic for data relations.

There is a large body of work on access control for privacy in collaborative systems, especially online social networks. Hu *et al.* develop multiparty access control, where they develop a social network model, a multiparty policy specification scheme and a mechanism to enforce policies to resolve multiparty privacy conflicts [46]. Carminati *et al.* study a semantic web based framework to manage access control in OSNs by generating semantic policies [20]. The social network operates according to agreed system-level policies. Fong [38] pioneered the application of relationship-based access control mechanisms to collaborative systems, which initiated different lines of research. The interoperability of relationship and role-based access control mechanisms is studied by Rizvi and Fong [77]. Mehregan and Fong [62] propose a policy negotiation mechanism for co-owned resources. These works provide feasible privacy resolution mechanisms for collaborative systems when policies are defined well. However, they require specification of policies for shared content offline, either by inference or with human expert involvement. Our work here, on the other hand, identifies the privacy norms that emerge in collaborative systems and makes them available to the users.

Norms have been studied in multiagent literature. Our earlier work investigated the idea of social norm emergence for OSNs [108]. However, it did not consider important aspects including prudential norms, aging of privacy decisions, or agent's autonomy in choosing to follow norms. Calikli *et al.* [19] employ a social identity map for relationships of users and a set of social identity conflict rules to learn the privacy norms for social networks. Mashayekhi *et al.* [61] study norm emergence in traffic domain, where agents enter and leave and no known network structure among them exists. Ajmeri *et al.* [3] study norm emergence factoring in the context of the agents, taking in the sanctions into account. Our work is orthogonal to these work in the sense that we investigate norms in privacy, where agents cannot be required to follow them.

Such *et al.* perform an extensive, empirical evaluation to understand the dynamics of privacy with co-owned content [96]. They indicate that various co-owner type relations as well as different type of handling of violations might exist. PRINOR could serve as a solution to the problems identified there, as by identifying norms and applying them as they see fit, the agents can avoid privacy violations to take place. Thuraisingham *et al.* [102] tackle the privacy-awareness in handling data that is collected for business and marketing purposes and discusses design issues in achieving privacy-aware data management frameworks. Since only a small amount of privacy policies are known for that domain, our approach could help identify social norms.

5.7.2 Limitations and Future Directions

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

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

In our approach, we have defined personal norms in a way that it captures the privacy understanding of users. These norms play a role in the emergence of social norms, since collaborative privacy decisions are the result of personal opinions. PRINOR captures this by regularly updating social norm bases to reflect the effect of personal choices on the society. However, we can intuitively say that users can also be influenced by emerging societal norms in such a way that their personal opinions about privacy changes. This would result in emerging personal norms, which can again affect the ever-changing societal norms. Implementing this notion can enable a normative privacy mechanism that would capture another layer of real-life behavior, resulting in even more accurate emergent norms. An extension like this can also provide the grounds for *explainability* of privacy behavior, both on personal and societal levels. To reach this, we have worked on information diffusion models to investigate how privacy behavior can spread among social network communities [72]. In this work, we created multi-agent simulation scenarios to demonstrate the dynamics of the spread and study the factors (e.g., trust) that influence the spread for privacy behavior. As a future direction, the results of these simulations can be integrated to a model that create a lifecycle for different types of norms, where each are affected by the decisions made by the others. With the addition of a model that captures this, PRINOR would yield more realistic results and can further help to assist OSN users in preserving privacy.



6

Conclusions and Future Directions

In this dissertation, we have tackled the challenges present in preserving privacy when a privacy decision affects more than an individual, which leads to collaborative privacy management. For our research, we have picked online social networks as the domain, which is a widely used one by users all around the world, where they can share content with other people. Since there are many open challenges in online social networks regarding privacy, it is a good testbed for privacy related research. Here, we mainly worked on providing a collaborative privacy mechanism for online social networks, with which software agents can automate privacy decisions for the users, while still considering human values and the social norms in their decisions.

Referring back to the three main research questions in Chapter 1, we list the conclusions we have reached with our research below.

Research Question 1: *What is an easy-to-use privacy resolution mechanism that enables collaborative privacy decisions in online social networks?* To tackle this research question, in this dissertation, we have developed an auction-based mechanism called PANO (explained in Chapter 2), which employs Clarke-Tax mechanism for its auctions and enables the participants to bid for privacy decisions. With PANO, each co-owner of a content receives an amount of units to spend for the auction, and they can bid on privacy actions to share the content publicly, to keep it private or to share it with a limited audience. Then the action with the highest total bid is applied, which becomes the collaborative decision of all co-owners. PANO contains features that prevents abuse of the system by a group of users who potentially can sway the decisions to their will, if those features were not available. We created multiagent simulations to evaluate PANO, with which we investigated success of the mechanism to reach correct privacy decisions and the user satisfaction of employing PANO as a collaborative privacy resolution mechanism. Our findings from this line of research are listed below.

- Auction-based mechanisms are useful for collaborative privacy decisions.
- PANO enables all users that co-own a content to have a say in the privacy decision, where each user is protected from abusive behavior and has an equal say in the outcome.

- PANO reaches successful privacy decisions, where oversharing is kept at a minimal level, which means privacy violations are rare occurrences.
- PANO also reduces undersharing significantly, which is beneficial for online social networks, where content sharing is one of the main interactions to build communities.
- PANO ensures satisfactory results for the majority of users for their privacy decisions. While 100% satisfaction is unobtainable due to conflicting privacy requirements of content co-owners, PANO manages to reach outcomes that are fitting with most of the users' policies, whilst privacy of every co-owner is still preserved.

Research Question 2: *How can we design software agents that assist online social network users on their privacy decisions?* PANO provides a mechanism where online social network users can express their privacy requirements and reach collaborative privacy decisions to share content. To assist users in these decisions, we have developed software agents called PANOLA, which are described in detail in Chapter 3. PANOLA learns from previous privacy decisions and user input and can represent the users in the auctions. PANOLA takes different privacy personas into consideration and treats each persona in a way that no user is left at disadvantage and can preserve its privacy. We list the results obtained from this body of work below.

- Learning users' privacy behavior in order to represent them is essential, since requiring user input for each collaborative privacy decision in online social networks is an impossible task, considering the tremendous amount of content shared within the widely used applications.
- PANOLA is an effective way to develop software agents that can represent online social network users for collaborative privacy decisions and preserve their privacy without requiring a significant amount of user input.
- PANOLA can assist users with different levels of knowledge about the privacy and with varying motivation levels to express their privacy concerns. The outcome of collaborative decisions show that each user is treated equally when PANOLA agents represent all the users where all achieve privacy preservation at similar levels.
- Since different users value individualism or conformism differently, software agents that represent users should take this into consideration. PANOLA offers this flexibility by considering these valuations in the learning process.
- PANOLA satisfies all the desirable properties that were defined in the literature for privacy management systems, which consist of *automation, fairness, concealment, protection before exposure, being easy-to-compute, robustness to cold start problem* and *dynamism*.

Research Question 3: *How do we identify human values and social norms in privacy decisions, and incorporate them in the privacy decisions?* To achieve our goals

for this research question, we first adopted Schwartz's value typology [81] for privacy decisions, with which we investigated the effects of online social network users' *openness to change* compared to their *conservationism*, and in another dimension, their *self-enhancement* values compared to their *self-transcendence*. The results show that values have effects in the privacy decisions and users with similar values would act similarly for collaborative privacy decisions. Building on top of this, we employed a norm-based privacy decision mechanism, called PRINOR, which identifies emergent privacy related social norms in an online social network, and prompts users to comply with these norms. The users also have the choice to not follow the social norms and decide with a collaborative privacy mechanism, such as PANO. We have conducted evaluations over multiagent simulations and showed that PRINOR can correctly identify privacy norms and their application results in privacy preserving outcomes for collaborative decisions. Below, we list our findings of our work related to privacy norms and the mechanism that makes use of these norms to reach privacy decisions.

- Human values play a role in the privacy decisions that online social network users make.
- People within a society can share similar privacy related values and constraints, resulting them to act in a similar manner in certain circumstances, which might form *social norms* for privacy decisions.
- Referring to social norms for privacy decisions offer an efficient way to reach collaborative decisions, where users can simply choose to follow the relevant norm while privacy preservation is still taken into consideration in the outcome.
- PRINOR identifies the emergent social norms both for the entire society and small groups within it, and makes use of them in privacy decisions successfully, even when most of the community in an online social network behaves heterogeneously.
- According to our case study, identified norms by PRINOR are in line with real-life social norms, which indicates it would be applicable to online social networks with a large amount of users, which would enable the mechanism to identify fine-grained norms for even more accurate collaborative privacy decisions.

While our proposed mechanisms can offer a solution to collaborative privacy management in online social networks, privacy still remains a hot topic on many domains with many challenges yet to be tackled. Therefore, there are many future directions that our research can take. PANO offers an efficient way to preserve privacy in online social networks. The auction mechanism at its core is Clarke-Tax mechanism, which is applicable in many different domains. Since there are many other domains that have open privacy challenges, such as Internet of Things, healthcare systems or cloud computing services, a generalization of PANO can offer a way that becomes a solution for privacy decisions of co-owned content. This can be achieved with a formalization of privacy actions, which can result in the same outcome regardless of the domain. With this approach, software agents that can perform in each domain

can be developed, which can also lead to easier data transfers or interactions between different applications.

Another future direction to our work would be the involvement of users in building privacy assistant software agents. To reach this, we are currently working on a user study where we offer a gamified version of collaborative privacy decisions with PANO to online social network users, and evaluate their behavior to understand the reasoning behind user choices. The results of similar user experiments can offer many ways to extend our work to make it more applicable in real-life cases. PANOLA agents can be updated to be more personalized in mimicking user behavior, which would improve both agents' performance and provide new implications to suggest new actions to the users. Another benefit would be to identify the social norms in the OSNs in an extensive way, where we can make use of the obtained data from the user study to have a better understanding of emergent norms.

With these future directions, the ultimate goal of preserving privacy in collaborative systems would be to have a mechanism that is accepted and employed by every application, which stays faithful to the principles of privacy to protect individuals' rights that were established over the years. While GDPR offered this to an extent for preserving privacy of the individuals when the information only pertains to a single entity, having a similar regulation which has an underlining collaborative privacy mechanism would help us to reach a solution where every online or offline application that contain co-owned data would be protected from privacy violations.



Bibliography

- [1] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015.
- [2] Nirav Ajmeri, Pradeep K Murukannaiah, Hui Guo, and Munindar P Singh. Arnor: Modeling social intelligence via norms to engineer privacy-aware personal agents. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 230–238, 2017.
- [3] Nirav Ajmeri, Hui Guo, Pradeep K Murukannaiah, and Munindar P Singh. 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)*, pages 22–34, 2018.
- [4] Nirav Ajmeri, Hui Guo, Pradeep K Murukannaiah, and Munindar P Singh. Elessar: Ethics in norm-aware agents. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 16–24, 2020.
- [5] Zeynep Akata, Dan Balliet, Maarten De Rijke, Frank Dignum, Virginia Dignum, Gusztai Eiben, Antske Fokkens, Davide Grossi, Koen Hindriks, Holger Hoos, Hayley Hung, Catholijn Jonker, Christof Monz, and Henry Prakken. A research agenda for hybrid intelligence: augmenting human intellect with collaborative, adaptive, responsible, and explainable artificial intelligence. *Computer*, 53(08): 18–28, 2020.
- [6] Cuneyt Akcora, Barbara Carminati, and Elena Ferrari. Privacy in social networks: How risky is your social graph? In *2012 IEEE 28th International Conference on Data Engineering*, pages 9–19. IEEE, 2012.
- [7] Cuneyt Gurcan Akcora, Barbara Carminati, and Elena Ferrari. User similarities on social networks. *Social Network Analysis and Mining*, 3(3):475–495, 2013.
- [8] Davide A. Albertini, Barbara Carminati, and Elena Ferrari. Privacy settings recommender for online social network. In *2016 IEEE 2nd International Conference on Collaboration and Internet Computing (CIC)*, pages 514–521, Nov 2016.

- [9] Natasha Alechina, Mehdi Dastani, and Brian Logan. Programming norm-aware agents. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '12*, pages 1057–1064, 2012.
- [10] Md. Zulfikar Alom, Barbara Carminati, and Elena Ferrari. Helping users managing context-based privacy preferences. In *IEEE International Conference on Services Computing (SCC)*, pages 100–107, 2019.
- [11] Georgia Bafoutsou and Gregoris Mentzas. Review and functional classification of collaborative systems. *International journal of information management*, 22(4):281–305, 2002.
- [12] Leila Bahri, Barbara Carminati, and Elena Ferrari. Decentralized privacy preserving services for online social networks. *Online Social Networks and Media*, 6:18–25, 2018. ISSN 2468-6964.
- [13] Adam Barth, Anupam Datta, John Mitchell, and Helen Nissenbaum. Privacy and contextual integrity: framework and applications. In *2006 IEEE Symposium on Security and Privacy (S P'06)*, pages 15 pp.–198, May 2006.
- [14] Adam Barth, John Mitchell, Anupam Datta, and Sharada Sundaram. Privacy and utility in business processes. In *20th IEEE Computer Security Foundations Symposium (CSF'07)*, pages 279–294, July 2007.
- [15] Andrew G. Barto and Sridhar Mahadevan. Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems*, 13(4):341–379, Oct 2003. ISSN 1573-7594.
- [16] David Basin, Matúš Harvan, Felix Klaedtke, and Eugen Zălinescu. Monopoly: Monitoring usage-control policies. In *Runtime Verification*, pages 360–364. Springer Berlin Heidelberg, 2012.
- [17] Elisa Bertino and Elena Ferrari. *Big Data Security and Privacy*, pages 425–439. Springer International Publishing, Cham, 2018.
- [18] Arthur A Bushkin and Samuel I Schaen. *The Privacy act of 1974: a reference manual for compliance*. System Development Corporation McLean, Va., 1976.
- [19] Gul Calikli, Mark Law, Arosha K. Bandara, Alessandra Russo, Luke Dickens, Blaine A. Price, Avelie Stuart, Mark Levine, and Bashar Nuseibeh. 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*, pages 47–56. ACM, 2016.
- [20] Barbara Carminati, Elena Ferrari, Raymond Heatherly, Murat Kantarcioglu, and Bhavani Thuraisingham. A semantic web based framework for social network access control. In *Proceedings of the 14th ACM symposium on Access control models and technologies*, pages 177–186. ACM, 2009.
- [21] Ann Cavoukian. Privacy by design: The 7 foundational principles. *Information and Privacy*, 2009.

- [22] Edward Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.
- [23] Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 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, page 1–13, New York, NY, USA, 2020. Association for Computing Machinery.
- [24] European Commission. General data protection regulation. Available at: <https://ec.europa.eu/commission/priorities/justice-and-fundamental-rights/data-protection/2018-reform-eu-data-protection-rules>, 2018.
- [25] European Commission. Directive 95/46/ec of the european parliament and of the council of 24 october 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. 1995.
- [26] Lorrie Cranor. *Web privacy with P3P*. O'Reilly Media, Inc., 2002.
- [27] Lorrie Faith Cranor. P3p: Making privacy policies more useful. *IEEE Security & Privacy*, 1(6):50–55, 2003.
- [28] Lorrie Faith Cranor, Joseph Reagle, and Mark S Ackerman. Beyond concern: Understanding net users' attitudes about online privacy. *The Internet upheaval: raising questions, seeking answers in communications policy*, pages 47–70, 2000.
- [29] Judith DeCew. Privacy. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring edition, 2018.
- [30] Francien Dechesne, Gennaro Di Tosto, Virginia Dignum, and Frank Dignum. No smoking here: values, norms and culture in multi-agent systems. *Artificial Intelligence and Law*, 21(1):79–107, Mar 2013. ISSN 1572-8382.
- [31] Carlos Diuk, Andre Cohen, and Michael L. Littman. An object-oriented representation for efficient reinforcement learning. In *Proceedings of the 25th International Conference on Machine Learning*, ICML '08, pages 240–247. ACM, 2008.
- [32] Jean-Baptist Du Prel, Gerhard Hommel, Bernd Röhrig, and Maria Blettner. Confidence interval or p-value?: part 4 of a series on evaluation of scientific publications. *Deutsches Ärzteblatt International*, 106(19):335, 2009.
- [33] Janna Lynn Dupree, Richard Devries, Daniel M Berry, and Edward Lank. Privacy personas: Clustering users via attitudes and behaviors toward security practices. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 5228–5239, 2016.

- [34] Eithan Ephrati and Jeffrey S. Rosenschein. The clarke tax as a consensus mechanism among automated agents. In *Proceedings of the Ninth National Conference on Artificial Intelligence - Volume 1, AAAI'91*, pages 173–178. AAAI Press, 1991.
- [35] Lujun Fang and Kristen LeFevre. Privacy wizards for social networking sites. In *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, pages 351–360. ACM, 2010.
- [36] R. L. Fogues, P. K. Murukannaiah, J. M. Such, and M. P. Singh. Sosharp: Recommending sharing policies in multiuser privacy scenarios. *IEEE Internet Computing*, 21(6):28–36, November 2017. ISSN 1089-7801.
- [37] Ricard L Fogues, Pradeep K Murukannaiah, Jose M Such, and Munindar P Singh. Sharing policies in multiuser privacy scenarios: Incorporating context, preferences, and arguments in decision making. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 24(1):1–29, 2017.
- [38] Philip W.L. Fong. Relationship-based access control: Protection model and policy language. In *Proceedings of the First ACM Conference on Data and Application Security and Privacy, CODASPY '11*, pages 191–202. ACM, 2011.
- [39] Organisation for Economic Co-operation and Development (OECD). Recommendation of the council concerning guidelines governing the protection of privacy and transborder flows of personal data. 1980.
- [40] Carrie Gates. Access control requirements for web 2.0 security and privacy. *IEEE Web*, 2(0):12–15, 2007.
- [41] Ken Gormley. One hundred years of privacy. *Wisconsin Law Review*, page 1335, 1992.
- [42] Barbara J Grosz. Collaborative systems (AAAI-94 presidential address). *AI magazine*, 17(2):67–67, 1996.
- [43] Seda Gürses and Claudia Diaz. Two tales of privacy in online social networks. *IEEE Security & Privacy*, 11(3):29–37, 2013.
- [44] Seda Gurses, Carmela Troncoso, and Claudia Diaz. Engineering privacy by design. *Computers, Privacy & Data Protection*, 2011. 25 pages.
- [45] Chris Haynes, Michael Luck, Peter McBurney, Samhar Mahmoud, Tomas Vitek, and Simon Miles. Engineering the emergence of norms: a review. *The Knowledge Engineering Review*, 32:e18, 2017.
- [46] Hongxin Hu, Gail-Joon Ahn, and Jan Jorgensen. Multiparty access control for online social networks: Model and mechanisms. *IEEE Transactions on Knowledge and Data Engineering*, 25(7):1614–1627, July 2013. ISSN 1041-4347.
- [47] Michael N Huhns and Munindar Paul Singh. *Readings in agents*. Morgan Kaufmann, 1998.

-
- [48] Panagiotis Ilia, Barbara Carminati, Elena Ferrari, Paraskevi Fragopoulou, and Sotiris Ioannidis. Sampac: Socially-aware collaborative multi-party access control. In *proceedings of the seventh ACM on conference on data and application security and privacy*, pages 71–82, 2017.
- [49] Simon Jones and Eamonn O’Neill. Feasibility of structural network clustering for group-based privacy control in social networks. In *proceedings of the sixth symposium on usable privacy and security*, pages 1–13, 2010.
- [50] Dilara Kekulluoglu, Nadin Kokciyan, and Pınar Yolum. Preserving privacy as social responsibility in online social networks. *ACM Transactions on Internet Technologies*, 18(4):42:1–42:22, April 2018. ISSN 1533-5399.
- [51] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Ashesh Rambachan. Algorithmic fairness. In *Aea papers and proceedings*, volume 108, pages 22–27, 2018.
- [52] Peter Klemperer, Yuan Liang, Michelle Mazurek, Manya Sleeper, Blase Ur, Lujo Bauer, Lorrie Faith Cranor, Nitin Gupta, and Michael Reiter. 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, pages 377–386. ACM, 2012.
- [53] Nadin Kökciyan and Pınar Yolum. Priguard: A semantic approach to detect privacy violations in online social networks. *IEEE Transactions on Knowledge and Data Engineering*, 28(10):2724–2737, 2016.
- [54] Nadin Kökciyan, Nefise Yaglikci, and Pınar Yolum. An argumentation approach for resolving privacy disputes in online social networks. *ACM Transactions on Internet Technologies*, 17(3):27:1–27:22, June 2017. ISSN 1533-5399.
- [55] Ponnurangam Kumaraguru and Lorrie Faith Cranor. Privacy indexes: A survey of westin’s studies. 2005. Available as *ISRI Technical Report CMU-ISRI-05-138*, 2005.
- [56] Airi Lampinen, Vilma Lehtinen, Asko Lehmuskallio, and Sakari Tamminen. 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, pages 3217–3226. ACM, 2011.
- [57] Marc Langheinrich. Privacy by design — principles of privacy-aware ubiquitous systems. In Gregory D. Abowd, Barry Brumitt, and Steven Shafer, editors, *Ubicomp 2001: Ubiquitous Computing*, pages 273–291, 2001.
- [58] Scott Lederer, Jason I Hong, Anind K Dey, and James A Landay. Personal privacy through understanding and action: five pitfalls for designers. *Personal and ubiquitous computing*, 8(6):440–454, 2004.
- [59] Jure Leskovec and Rok Sosič. Snap: A general-purpose network analysis and graph-mining library. *ACM Trans. Intell. Syst. Technol.*, 8(1), July 2016. ISSN 2157-6904.

- [60] I Martinez-Moyano. Exploring the dynamics of collaboration in interorganizational settings. *Creating a culture of collaboration: The International Association of Facilitators handbook*, 4:69, 2006.
- [61] Mehdi Mashayekhi, Hongying Du, George F List, and Munindar P Singh. Silk: A simulation study of regulating open normative multiagent systems. In *Proceedings of the International Joint Conference on AI (IJCAI)*, pages 373–379, 2016.
- [62] Pooya Mehregan and Philip W.L. Fong. Policy negotiation for co-owned resources in relationship-based access control. In *Proceedings of the 21st ACM on Symposium on Access Control Models and Technologies, SACMAT '16*, page 125–136. Association for Computing Machinery, 2016.
- [63] Yavuz Mester, Nadin Kökciyan, and Pınar Yolum. *Negotiating Privacy Constraints in Online Social Networks*, pages 112–129. Springer International Publishing, Cham, 2015.
- [64] Gaurav Misra and Jose M Such. Pacman: Personal agent for access control in social media. *IEEE Internet Computing*, 21(6):18–26, 2017.
- [65] Javier Morales, Maite Lopez-Sanchez, Juan A. Rodriguez-Aguilar, Michael Wooldridge, and Wamberto Vasconcelos. Automated synthesis of normative systems. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems, AAMAS '13*, pages 483–490, 2013.
- [66] Francesca Mosca. Value-aligned and explainable agents for collective decision making: Privacy application. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 2199–2200, 2020.
- [67] Francesca Mosca and Jose Such. Elvira: An explainable agent for value and utility-driven multiuser privacy. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2021.
- [68] Francesca Mosca and Jose Such. An explainable assistant for multiuser privacy. *Autonomous Agents and Multi-Agent Systems*, 36(1):1–45, 2022.
- [69] Francesca Mosca, Jose M Such, and Peter John McBurney. Value-driven collaborative privacy decision making. In *Proceedings of AAAI Spring Symposium 2019: PAL: Privacy-Enhancing Artificial Intelligence and Language Technologies*, 2018.
- [70] Francesca Mosca, Jose M Such, and Peter McBurney. Towards a value-driven explainable agent for collective privacy. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1937–1939, 2020.
- [71] Deirdre K. Mulligan and Colin Koopman. Theorizing privacy’s contestability: A multi-dimensional analytic of privacy. In *Proceedings of Special Workshop on Information Privacy*, 2015.

- [72] Albert Mwanjesa, Onuralp Ulusoy, and Pinar Yolum. Dipp: Diffusion of privacy preferences in online social networks. In *Advances in Social Simulation*, pages 29–40. Springer, 2022.
- [73] Judith S Olson, Jonathan Grudin, and Eric Horvitz. A study of preferences for sharing and privacy. In *CHI'05 extended abstracts on Human factors in computing systems*, pages 1985–1988, 2005.
- [74] Federica Paci, Anna Squicciarini, and Nicola Zannone. Survey on access control for community-centered collaborative systems. *ACM Computing Surveys*, 51(1):6:1–6:38, January 2018. ISSN 0360-0300.
- [75] William Prosser. Privacy. *California Law Review*, 48(3):41, 1960.
- [76] Sarah Rajtmajer, Anna Squicciarini, Jose M Such, Justin Semonsen, and Andrew Belmonte. An ultimatum game model for the evolution of privacy in jointly managed content. In *International Conference on Decision and Game Theory for Security*, pages 112–130. Springer, 2017.
- [77] Syed Zain R. Rizvi and Philip W.L. Fong. Interoperability of relationship- and role-based access control. In *Proceedings of the Sixth ACM Conference on Data and Application Security and Privacy*, CODASPY '16, page 231–242, 2016.
- [78] R. S. Sandhu, E. J. Coyne, H. L. Feinstein, and C. E. Youman. Role-based access control models. *Computer*, 29(2):38–47, Feb 1996. ISSN 0018-9162.
- [79] Ravi S. Sandhu, Edward J. Coyne, Hal L. Feinstein, and Charles E. Youman. Role-based access control models. *IEEE Computer*, 29(2):38–47, 1996.
- [80] Bastin Tony Roy Savarimuthu and Stephen Cranefield. Norm creation, spreading and emergence: A survey of simulation models of norms in multi-agent systems. *Multiagent Grid Syst.*, 7(1):21–54, January 2011. ISSN 1574-1702.
- [81] Shalom H Schwartz. An overview of the schwartz theory of basic values. *Online readings in Psychology and Culture*, 2(1):11, 2012.
- [82] Sandip Sen and Stéphane Airiau. Emergence of norms through social learning. In *Proceedings of the International Joint Conference on AI (IJCAI)*, volume 1507, page 1512, 2007.
- [83] Bikash Chandra Singh, Barbara Carminati, and Elena Ferrari. Privacy-aware personal data storage (p-pds): Learning how to protect user privacy from external applications. *IEEE Transactions on Dependable and Secure Computing*, 18(2):889–903, 2019.
- [84] Daniel J Solove. A taxonomy of privacy. In *University of Pennsylvania Law Review*, volume 154, page 477–560, 2006.
- [85] Daniel J Solove. *Understanding privacy*. Harvard University Press, Cambridge, May, 2008.

- [86] Sarah Spiekermann. The challenges of privacy by design. *Communications of the ACM*, 55(7):38–40, 2012.
- [87] Sarah Spiekermann and Lorrie Faith Cranor. Engineering privacy. *IEEE Trans. Softw. Eng.*, 35(1):67–82, January 2009. ISSN 0098-5589.
- [88] Anna C. Squicciarini, Mohamed Shehab, and Federica Paci. Collective privacy management in social networks. In *Proceedings of the 18th International Conference on World Wide Web, WWW '09*, pages 521–530. ACM, 2009.
- [89] Anna C. Squicciarini, Smitha Sundareswaran, Dan Lin, and Josh Wede. A3p: Adaptive policy prediction for shared images over popular content sharing sites. In *Proceedings of the 22nd ACM conference on Hypertext and hypermedia*, pages 261–270, 2011.
- [90] Anna C Squicciarini, Federica Paci, and Smitha Sundareswaran. Prima: a comprehensive approach to privacy protection in social network sites. *Annals of Telecommunications-Annales des Télécommunications*, 69(1):21–36, 2014.
- [91] Anna C. Squicciarini, Dan Lin, Smitha Sundareswaran, and Joshua Wede. Privacy policy inference of user-uploaded images on content sharing sites. *IEEE Transactions on Knowledge and Data Engineering*, 27(1):193–206, Jan 2015. ISSN 1041-4347.
- [92] Anna C. Squicciarini, Cornelia Caragea, and Rahul Balakavi. Toward automated online photo privacy. *ACM Transactions on the Web*, 11(1):2:1–2:29, April 2017. ISSN 1559-1131.
- [93] Jose M. Such and Natalia Criado. Resolving multi-party privacy conflicts in social media. *IEEE Transactions on Knowledge and Data Engineering*, 28(7): 1851–1863, July 2016. ISSN 1041-4347.
- [94] Jose M. Such and Natalia Criado. Multiparty privacy in social media. *Communications of the ACM*, 61(8):74–81, July 2018. ISSN 0001-0782.
- [95] Jose M. Such and Michael Rovatsos. Privacy policy negotiation in social media. *ACM Transactions on Autonomous and Adaptive Systems*, 11(1):4:1–4:29, February 2016. ISSN 1556-4665.
- [96] Jose M Such, Joel Porter, Sören Preibusch, and Adam Joinson. Photo privacy conflicts in social media: A large-scale empirical study. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 3821–3832. ACM, 2017.
- [97] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [98] Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *In Proceedings of the Tenth International Conference on Machine Learning*, pages 330–337. Morgan Kaufmann, 1993.

-
- [99] Pang-Ning Tan, Michael Steinbach, and Vipin Kumar. *Introduction to Data Mining, (First Edition)*. Addison-Wesley Longman Publishing Co., Inc., 2005.
- [100] Roshan K Thomas. Team-based access control (tmac) a primitive for applying role-based access controls in collaborative environments. In *Proceedings of the second ACM workshop on Role-based access control*, pages 13–19, 1997.
- [101] Roshan K Thomas and Ravi S Sandhu. Task-based authorization controls (tbac): A family of models for active and enterprise-oriented authorization management. In *Database security XI*, pages 166–181. Springer, 1998.
- [102] Bhavani Thuraisingham, Murat Kantarcioglu, Elisa Bertino, Jonathan Z. Bakhdash, and Maribel Fernandez. Towards a privacy-aware quantified self data management framework. In *Proceedings of the 23rd ACM on Symposium on Access Control Models and Technologies*, SACMAT '18, page 173–184, 2018.
- [103] William Tolone, Gail-Joon Ahn, Tanusree Pai, and Seng-Phil Hong. Access control in collaborative systems. *ACM Computing Surveys (CSUR)*, 37(1):29–41, 2005.
- [104] Ashwini Tonge, Cornelia Caragea, and Anna Squicciarini. Uncovering scene context for predicting privacy of online shared images. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32-1, 2018.
- [105] Raimo Tuomela. *The Importance of Us: A Philosophical Study of Basic Social Norms*. Stanford University Press, 01 1995.
- [106] Onuralp Ulusoy. Collaborative privacy management in online social networks. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1788–1790, 2018.
- [107] Onuralp Ulusoy and Pinar Yolum. Pano: Privacy auctioning for online social networks. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '18, pages 2103–2105, 2018.
- [108] Onuralp Ulusoy and Pinar Yolum. Emergent privacy norms for collaborative systems. In *PRIMA 2019: Principles and Practice of Multi-Agent Systems*, pages 514–522, Cham, 2019. Springer International Publishing.
- [109] Onuralp Ulusoy and Pinar Yolum. Privacy norms in online social networks. In *Benelux Conference on Artificial Intelligence*, 2019.
- [110] Onuralp Ulusoy and Pinar Yolum. Agents for preserving privacy: Learning and decision making collaboratively. In Nick Bassiliades, Georgios Chalkiadakis, and Dave de Jonge, editors, *Multi-Agent Systems and Agreement Technologies*, pages 116–131. Springer International Publishing, 2020.
- [111] Onuralp Ulusoy and Pinar Yolum. Norm-based access control. In *Proceedings of the 25th ACM Symposium on Access Control Models and Technologies*, SACMAT '20, page 35–46, New York, NY, USA, 2020. Association for Computing Machinery.

- [112] Onuralp Ulusoy and Pinar Yolum. Collaborative privacy management with auctioning mechanisms. In *Advances in Automated Negotiations*, pages 45–62, Singapore, 2020. Springer Singapore.
- [113] Onuralp Ulusoy and Pinar Yolum. Panola: A personal assistant for supporting users in preserving privacy. *ACM Transactions on Internet Technologies*, 22(1), sep 2021. ISSN 1533-5399.
- [114] M. Vanetti, E. Binaghi, E. Ferrari, B. Carminati, and M. Carullo. A system to filter unwanted messages from osn user walls. *IEEE Transactions on Knowledge and Data Engineering*, 25(2):285–297, Feb 2013. ISSN 1041-4347.
- [115] Samuel Warren and Louis Brandeis. The right to privacy. In *Killing the Messenger*, pages 1–21. Columbia University Press, 1989.
- [116] Alan F Westin. Privacy and freedom. *Washington and Lee Law Review*, 25(1):166, 1968.
- [117] Piotr A. Woźniak, Edward J. Gorzelańczyk, and Janusz A. Murakowski. Two components of long-term memory. *Acta neurobiologiae experimentalis*, 55(4):301–305, 1995. ISSN 0065-1400.
- [118] Matthew D. Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 818–833, Cham, 2014.
- [119] Sergej Zerr, Stefan Siersdorfer, and Jonathon Hare. Picalert! a system for privacy-aware image classification and retrieval. In *Proceedings of the 21st ACM International Conference on Information and Knowledge Management, CIKM ’12*, page 2710–2712. Association for Computing Machinery, 2012.
- [120] Haoti Zhong, Anna C. Squicciarini, and David Miller. Toward automated multiparty privacy conflict detection. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM ’18*, pages 1811–1814, New York, NY, USA, 2018. ACM.
- [121] Haibin Zhu. Role mechanisms in collaborative systems. *International Journal of Production Research*, 44(1):181–193, 2006.
- [122] Michael Zimmer and Anthony Hoffman. Privacy, context, and oversharing: Reputational challenges in a web 2.0 world. *The reputation society: How online opinions are reshaping the offline world*, page 175, 2012.



Summary

Privacy is the right of individuals to keep personal information to themselves or restrict the information to those that they are comfortable sharing with. 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 that are affected by the information have different, possibly conflicting privacy constraints.

In this dissertation, we investigate collaborative privacy mechanisms that can provide a fair resolution that also considers equity for preserving the privacy of all parties in question. We mainly focus on the online social networks domain, since it is widely used and contains pieces of content that can affect privacy of more than just one individual. We propose an auction-based mechanism, where each related party bids for a privacy preference for a collaborative privacy decision. The mechanism aims to be robust, immune to abuse and fair to all parties, so that each individual has a similar amount of effect in privacy decisions. We also prevent undersharing as well as oversharing, to ensure each piece of content that does not depict a privacy violation is shared in the system. We employ software agents to assist the users to participate in this mechanism, which is capable of learning privacy preference of all users with different levels of knowledge and motivation. Privacy assistant agents provide privacy preserving resolutions to conflicts where the users do not have to spend time and effort and can use the online social networks in a similar way that are provided by the widely used applications. Furthermore, we investigate the effect of human values in privacy decisions, in order to have an understanding of the thinking process behind online social network users. This leads us to social norms, which emerge from a group of users behaving in a similar way in similar circumstances. We propose a mechanism to identify these norms and incorporate them in collaborative privacy decisions, where users have the choice to simply follow the norms without the need of a complex decision system, and can still preserve privacy while preventing significant privacy violations. We evaluate our work with multi-agent simulations and

case studies and report the results in terms of success in preserving privacy, equity and usability. We show that our work offers an easy-to-use solution for preventing privacy violations, where each user's privacy concerns are taken into account, regardless of their knowledge about privacy or their motivation to be a part of the process.



Samenvatting

Privacy is het recht van individuen om persoonlijke informatie voor zichzelf te houden of de informatie te beperken tot degenen met wie ze zich op hun gemak voelen. Wanneer individuen online systemen gebruiken, moeten ze het recht krijgen om te beslissen welke informatie ze willen delen en wat ze privé willen houden. Wanneer een stuk informatie slechts betrekking heeft op één persoon, is het mogelijk om de privacy te behouden door de gebruiker de juiste toegangsopties te bieden. Wanneer een stukje informatie echter betrekking heeft op meerdere individuen, zoals een foto van een groep vrienden of een gezamenlijk bewerkt document, is het de uitdaging om te beslissen hoe deze informatie te delen en met wie. Het probleem wordt moeilijker wanneer de personen die door de informatie worden geraakt, verschillende, mogelijk tegenstrijdige privacybeperkingen hebben.

In dit proefschrift onderzoeken we collaboratieve privacymechanismen die een eerlijke oplossing kunnen bieden die ook rekening houdt met rechtvaardigheid voor het behoud van de privacy van alle betrokken partijen. We richten ons voornamelijk op het domein van online sociale netwerken, omdat deze veel worden gebruikt en stukjes inhoud bevatten die de privacy van meer dan één persoon kunnen beïnvloeden. We stellen een op veilingen gebaseerd mechanisme voor, waarbij elke verbonden partij biedt op een privacyvoorkeur voor een gezamenlijke privacybeslissing. Het mechanisme moet robuust zijn, immuun voor misbruik en eerlijk zijn voor alle partijen, zodat elk individu de zelfde invloed heeft bij privacybeslissingen. We voorkomen ook zowel het te weinig als het te veel delen, om ervoor te zorgen dat elk stukje inhoud dat geen privacyschending bevat, in het systeem wordt gedeeld. We gebruiken softwareagenten om de gebruikers te helpen deel te nemen aan dit mechanisme, dat in staat is om de privacyvoorkeuren van alle gebruikers met verschillende kennisniveaus en motivatie te leren. Privacyassistent-agenten bieden privacybeschermende oplossingen voor conflicten waarbij de gebruikers geen tijd en moeite hoeven te besteden en de online sociale netwerken op een vergelijkbare manier kunnen gebruiken als met de veelgebruikte applicaties. Verder onderzoeken we het effect van menselijke waarden op privacybeslissingen, om inzicht te krijgen in het denkproces achter gebruikers van online sociale netwerken. Dit leidt ons naar sociale normen, die voortkomen uit een groep gebruikers die zich in vergelijkbare omstandigheden op dezelfde manier gedraagt. We stellen een mechanisme voor om deze normen te identificeren en op

te nemen in gezamenlijke privacybeslissingen, waarbij gebruikers de keuze hebben om simpelweg de normen te volgen zonder een complex beslissingssysteem te hoeven gebruiken, en toch de privacy kunnen behouden en tegelijkertijd aanzienlijke privacy-schendingen kunnen voorkomen. We evalueren ons werk met multi-agent simulaties en case studies en rapporteren de resultaten in termen van succes met betrekking tot het behoud van privacy, gelijkheid en bruikbaarheid. We laten zien dat ons werk een gebruiksvriendelijke oplossing biedt voor het voorkomen van privacy-schendingen, waarbij rekening wordt gehouden met de privacyzorgen van elke gebruiker, ongeacht hun kennis van privacy of hun motivatie om deel uit te maken van het proces.



Acknowledgements

It has been a long journey from the beginning of my PhD to reach this point to publish a condensed version of my work. Of course it would not have been possible without the support of people that helped me along the way. First and foremost, I would like to express my deepest gratitude to my supervisor Pinar Yolum, who has always been guiding and supporting me throughout the years to achieve this. You immensely helped me grow to be a better researcher, and without you I certainly would not have achieved the things I did. Working with you has been a huge pleasure and I hope we will be continuing to do so in the coming years.

I would like to thank my co-promotor Tim Baarslag for his valuable feedback for my research, which helped me to improve this dissertation in a way I would not have done by myself. I would again like to thank him, Pinar and Ilaria Liccardi for being a part of my privacy user experiment project for the last year, and I hope we will be able to make it reach to its conclusion in the near future. I would also like to thank the members of my assessment committee Frank Dignum, Catholijn Jonker, Albert Ali Salah, Anna Squicciarini and Jose Such for evaluating this dissertation and providing me with valuable feedback. With your constructive comments I was able to make the final touches to improve my thesis.

Of course, reaching the end of this long journey would not have been possible without friends. First, I would like to thank Can, who has been a great friend, a colleague and a roommate for a part of my PhD. Sharing a similar experience of moving to the Netherlands with you helped me a lot for dealing with obstacles and not struggle with my new life here. Thank you very much to Davide, who has not only been a good friend to spend time together, but also a very supportive colleague. Your experiences on your own PhD journey has been a guide to me and it has been a pleasure to share a part of that experience with you. Thank you Jan for being a good friend throughout my PhD, and also thanks for helping me with the Dutch summary translation of this dissertation. I would like to thank all my PhD colleagues, who helped me to have a great experience over the years of my time in the department. You have always been good supportive friends that helped me keep my spirits high while working towards reaching this achievement in my life. I will try to name you all and I am sorry if I forget any names, but I very much appreciate you all being a part of my life. Thank you very much Daphne, Maarten, Danping, Jurian, Remi, Marcel,

Samaneh, Gönül, Isaac, Mijke, Bilge, Alexander, Bas, Federica, Anouk, Vinicius, Metehan, Michiel, Emre E., Duygu, Emre O., Annet, Maksim, Changxi and again many others that I could not name here. I wish the ones who are still on their PhD journey a smooth and successful path to reach where I am now, and the ones who already finished a very happy life ahead.

I would also like to thank all the professors, lecturers and post-docs I have met throughout my PhD. Every brief or detailed conversation we had had a contribution to my research and I am thankful for sharing the same work environment with you. Thank you to the employees of our department and Utrecht University, who supported me in many occasions to answer my questions and solve the issues I have faced. Thank you Albert for choosing me as your MSc supervisor. It was a pleasure to work with you and be a part of your achievements on your way to your Masters degree. I wish you a very successful career ahead. I would like to thank my friends outside academia, my family members and friends I have made in conferences and summer schools. All of you have made my life happier and have a great contribution that kept me motivated to reach where I am now.

Finally, a very special thanks to my parents Erdem and Fatma, for always being there for me. You have endured all my struggles, stress and concerns with me for my entire life, and it would have been impossible to get over them to have this achievement if not for you. Without you, I would not have been the person I am and I am eternally thankful to both of you.

Onuralp Ulusoy
Utrecht, 12/10/2022



Curriculum Vitae

Education

- 2018 - 2022 **PhD in Computer Science**
Utrecht University, Dept. of Information and Computing Science
“Privacy in Collaborative Systems”
Promotor: Prof. Dr. P. Yolum Birbil
Co-promotor: Dr. T. Baarslag
- 2009 - 2012 **MSc. Degree in Computer Science**
Istanbul Technical University, Dept. of Computer Science
“Distributed Team Formation for Robot Soccer”
Advisor: Assoc. Prof. Dr. S. Sariel
- 2005 - 2009 **BSc. Degree in Computer Science**
Trakya University, Dept. of Computer Science
“A natural language dialogue system for autistic children”
Advisor: Prof. Dr. Y. Kılıçaslan

Academic Activities

- 2018 - 2022 **Teaching Assistant**
Utrecht University, Dept. of Information and Computing Science
Courses: Methods in AI Research (4 terms), Intelligente Systemen (3 terms), Introductie Informatica (1 term)
- 2020 - 2021 **Student Supervision**
Utrecht University, Dept. of Information and Computing Science
MSc: DIPP: Information Diffusion for Privacy in Multi-agent Systems (2020-2021)

Involvement in Scientific Events

Program Committee

- EUMAS 2022 The 19th European Conference on Multi-Agent Systems
- AI4P 2020 Workshop on AI for Privacy at ECAI 2020

Reviewer for International Journals

- Springer Computer Science
- Nature

Reviewer for International Conferences

- EUMAS 2022 The 19th European Conference on Multi-Agent Systems
- AI4P 2020 Workshop on AI for Privacy at ECAI 2020
- EUMAS 2020 The 17th European Conference on Multi-Agent Systems
- PAAMS 2020 The 18th International Conference on Practical Applications of Agents and Multi-Agent Systems



SIKS Dissertation Series

-
- 2011 01 Botond Cseke (RUN), Variational Algorithms for Bayesian Inference in Latent Gaussian Models
 - 02 Nick Tinnemeier (UU), Organizing Agent Organizations. Syntax and Operational Semantics of an Organization-Oriented Programming Language
 - 03 Jan Martijn van der Werf (TUE), Compositional Design and Verification of Component-Based Information Systems
 - 04 Hado van Hasselt (UU), Insights in Reinforcement Learning; Formal analysis and empirical evaluation of temporal-difference
 - 05 Bas van der Raadt (VU), Enterprise Architecture Coming of Age - Increasing the Performance of an Emerging Discipline.
 - 06 Yiwen Wang (TUE), Semantically-Enhanced Recommendations in Cultural Heritage
 - 07 Yujia Cao (UT), Multimodal Information Presentation for High Load Human Comp. Interaction
 - 08 Nieske Vergunst (UU), BDI-based Generation of Robust Task-Oriented Dialogues
 - 09 Tim de Jong (OU), Contextualised Mobile Media for Learning
 - 10 Bart Bogaert (UvT), Cloud Content Contention
 - 11 Dhaval Vyas (UT), Designing for Awareness: An Experience-focused HCI Perspective
 - 12 Carmen Bratosin (TUE), Grid Architecture for Distributed Process Mining
 - 13 Xiaoyu Mao (UvT), Airport under Control Multiagent Scheduling for Airport Ground Handling
 - 14 Milan Lovric (EUR), Behavioral Finance and Agent-Based Artificial Markets
 - 15 Marijn Koolen (UvA), The Meaning of Structure: the Value of Link Evidence for Information Retrieval
 - 16 Maarten Schadd (UM), Selective Search in Games of Different Complexity
 - 17 Jiyin He (UVA), Exploring Topic Structure: Coherence, Diversity and Relatedness
 - 18 Mark Ponsen (UM), Strategic Decision-Making in complex games
 - 19 Ellen Rusman (OU), The Mind's Eye on Personal Profiles
 - 20 Qing Gu (VU), Guiding service-oriented software engineering - A view-based approach
 - 21 Linda Terlouw (TUD), Modularization and Specification of Service-Oriented Systems
 - 22 Junte Zhang (UVA), System Evaluation of Archival Description and Access
 - 23 Wouter Weerkamp (UVA), Finding People and their Utterances in Social Media
 - 24 Herwin van Welbergen (UT), Behavior Generation for Interpersonal Coordination with Virtual Humans On Specifying, Scheduling and Realizing Multimodal Virtual Human Behavior
 - 25 Syed Waqar ul Qounain Jaffry (VU), Analysis and Validation of Models for Trust Dynamics
 - 26 Matthijs Aart Pontier (VU), Virtual Agents for Human Communication - Emotion Regulation and Involvement-Distance Trade-Offs in Embodied Conversational Agents and Robots
 - 27 Aniel Bhulai (VU), Dynamic website optimization through autonomous management of design patterns
 - 28 Rianne Kaptein (UVA), Effective Focused Retrieval by Exploiting Query Context and Document Structure
 - 29 Faisal Kamiran (TUE), Discrimination-aware Classification
 - 30 Egon van den Broek (UT), Affective Signal Processing: Unraveling the mystery of emotions
 - 31 Ludo Waltman (EUR), Computational and Game-Theoretic Approaches for Modeling Bounded Rationality
 - 32 Nees-Jan van Eck (EUR), Methodological Advances in Bibliometric Mapping of Science
 - 33 Tom van der Weide (UU), Arguing to Motivate Decisions
 - 34 Paolo Turrini (UU), Strategic Reasoning in Interdependence: Logical and Game-theoretical Investigations
 - 35 Maaïke Harbers (UU), Explaining Agent Behavior in Virtual Training
 - 36 Erik van der Spek (UU), Experiments in serious game design: a cognitive approach

- 37 Adriana Burlutiu (RUN), Machine Learning for Pairwise Data, Applications for Preference Learning and Supervised Network Inference
 - 38 Nyree Lemmens (UM), Bee-inspired Distributed Optimization
 - 39 Joost Westra (UU), Organizing Adaptation using Agents in Serious Games
 - 40 Viktor Clerc (VU), Architectural Knowledge Management in Global Software Development
 - 41 Luan Ibraimi (UT), Cryptographically Enforced Distributed Data Access Control
 - 42 Michal Sindlar (UU), Explaining Behavior through Mental State Attribution
 - 43 Henk van der Schuur (UU), Process Improvement through Software Operation Knowledge
 - 44 Boris Reuderink (UT), Robust Brain-Computer Interfaces
 - 45 Herman Stehouwer (UvT), Statistical Language Models for Alternative Sequence Selection
 - 46 Beibei Hu (TUD), Towards Contextualized Information Delivery: A Rule-based Architecture for the Domain of Mobile Police Work
 - 47 Azizi Bin Ab Aziz (VU), Exploring Computational Models for Intelligent Support of Persons with Depression
 - 48 Mark Ter Maat (UT), Response Selection and Turn-taking for a Sensitive Artificial Listening Agent
 - 49 Andreea Niculescu (UT), Conversational interfaces for task-oriented spoken dialogues: design aspects influencing interaction quality
-
- 2012 01 Terry Kakeeto (UvT), Relationship Marketing for SMEs in Uganda
 - 02 Muhammad Umair (VU), Adaptivity, emotion, and Rationality in Human and Ambient Agent Models
 - 03 Adam Vanya (VU), Supporting Architecture Evolution by Mining Software Repositories
 - 04 Jurriaan Souer (UU), Development of Content Management System-based Web Applications
 - 05 Marijn Plomp (UU), Maturing Interorganisational Information Systems
 - 06 Wolfgang Reinhardt (OU), Awareness Support for Knowledge Workers in Research Networks
 - 07 Rianne van Lambalgen (VU), When the Going Gets Tough: Exploring Agent-based Models of Human Performance under Demanding Conditions
 - 08 Gerben de Vries (UVA), Kernel Methods for Vessel Trajectories
 - 09 Ricardo Neisse (UT), Trust and Privacy Management Support for Context-Aware Service Platforms
 - 10 David Smits (TUE), Towards a Generic Distributed Adaptive Hypermedia Environment
 - 11 J.C.B. Rantham Prabhakara (TUE), Process Mining in the Large: Preprocessing, Discovery, and Diagnostics
 - 12 Kees van der Sluijs (TUE), Model Driven Design and Data Integration in Semantic Web Information Systems
 - 13 Suleman Shahid (UvT), Fun and Face: Exploring non-verbal expressions of emotion during playful interactions
 - 14 Evgeny Knutov (TUE), Generic Adapt. Framework for Unifying Adaptive Web-based Systems
 - 15 Natalie van der Wal (VU), Social Agents. Agent-Based Modelling of Integrated Internal and Social Dynamics of Cognitive and Affective Processes.
 - 16 Fiemke Both (VU), Helping people by understanding them - Ambient Agents supporting task execution and depression treatment
 - 17 Amal Elgammal (UvT), Towards a Comprehensive Framework for Business Process Compliance
 - 18 Eltjo Poort (VU), Improving Solution Architecting Practices
 - 19 Helen Schonenberg (TUE), What's Next? Operational Support for Business Process Execution
 - 20 Ali Bahramisharif (RUN), Covert Visual Spatial Attention, a Robust Paradigm for Brain-Computer Interfacing
 - 21 Roberto Cornacchia (TUD), Querying Sparse Matrices for Information Retrieval
 - 22 Thijs Vis (UvT), Intelligence, politie en veiligheidsdienst: verenigbare grootheden?
 - 23 Christian Muehl (UT), Toward Affective Brain-Computer Interfaces: Exploring the Neurophysiology of Affect during Human Media Interaction
 - 24 Laurens van der Werff (UT), Evaluation of Noisy Transcripts for Spoken Document Retrieval
 - 25 Silja Eckartz (UT), Managing the Business Case Development in Inter-Organizational IT Projects: A Methodology and its Application
 - 26 Emile de Maat (UVA), Making Sense of Legal Text
 - 27 Hayrettin Gurkok (UT), Mind the Sheep! User Experience Evaluation & Brain-Computer Interface Games
 - 28 Nancy Pascall (UvT), Engendering Technology Empowering Women
 - 29 Almer Tigelaar (UT), Peer-to-Peer Information Retrieval
 - 30 Alina Pommeranz (TUD), Designing Human-Centered Systems for Reflective Decision Making
 - 31 Emily Bagarukayo (RUN), A Learning by Construction Approach for Higher Order Cognitive Skills Improvement, Building Capacity and Infrastructure
 - 32 Wietske Visser (TUD), Qualitative multi-criteria preference representation and reasoning
 - 33 Rory Sie (OUN), Coalitions in Cooperation Networks (COCOON)
 - 34 Pavol Jancura (RUN), Evolutionary analysis in PPI networks and applications
 - 35 Evert Haasdijk (VU), Never Too Old To Learn - On-line Evolution of Controllers in Swarm- and Modular Robotics
 - 36 Denis Ssebugwawo (RUN), Analysis and Evaluation of Collaborative Modeling Processes

-
- 37 Agnes Nakakawa (RUN), A Collaboration Process for Enterprise Architecture Creation
38 Selmar Smit (VU), Parameter Tuning and Scientific Testing in Evolutionary Algorithms
39 Hassan Fatemi (UT), Risk-aware design of value and coordination networks
40 Agus Gunawan (UvT), Information Access for SMEs in Indonesia
41 Sebastian Kelle (OU), Game Design Patterns for Learning
42 Dominique Verpoorten (OU), Reflection Amplifiers in self-regulated Learning
43 Withdrawn
44 Anna Tordai (VU), On Combining Alignment Techniques
45 Benedikt Kratz (UvT), A Model and Language for Business-aware Transactions
46 Simon Carter (UVA), Exploration and Exploitation of Multilingual Data for Statistical Machine Translation
47 Manos Tsagkias (UVA), Mining Social Media: Tracking Content and Predicting Behavior
48 Jorn Bakker (TUE), Handling Abrupt Changes in Evolving Time-series Data
49 Michael Kaisers (UM), Learning against Learning - Evolutionary dynamics of reinforcement learning algorithms in strategic interactions
50 Steven van Kervel (TUD), Ontology driven Enterprise Information Systems Engineering
51 Jeroen de Jong (TUD), Heuristics in Dynamic Scheduling; a practical framework with a case study in elevator dispatching
-
- 2013 01 Viorel Milea (EUR), News Analytics for Financial Decision Support
02 Erietta Liarou (CWI), MonetDB/DataCell: Leveraging the Column-store Database Technology for Efficient and Scalable Stream Processing
03 Szymon Klarman (VU), Reasoning with Contexts in Description Logics
04 Chetan Yadati (TUD), Coordinating autonomous planning and scheduling
05 Dulce Pumareja (UT), Groupware Requirements Evolutions Patterns
06 Romulo Goncalves (CWI), The Data Cyclotron: Juggling Data and Queries for a Data Warehouse Audience
07 Giel van Lankveld (UvT), Quantifying Individual Player Differences
08 Robbert-Jan Merk (VU), Making enemies: cognitive modeling for opponent agents in fighter pilot simulators
09 Fabio Gori (RUN), Metagenomic Data Analysis: Computational Methods and Applications
10 Jeewanie Jayasinghe Arachchige (UvT), A Unified Modeling Framework for Service Design.
11 Evangelos Pournaras (TUD), Multi-level Reconfigurable Self-organization in Overlay Services
12 Marian Razavian (VU), Knowledge-driven Migration to Services
13 Mohammad Safiri (UT), Service Tailoring: User-centric creation of integrated IT-based home-care services to support independent living of elderly
14 Jafar Tanha (UVA), Ensemble Approaches to Semi-Supervised Learning
15 Daniel Hennes (UM), Multiagent Learning - Dynamic Games and Applications
16 Eric Kok (UU), Exploring the practical benefits of argumentation in multi-agent deliberation
17 Koen Kok (VU), The PowerMatcher: Smart Coordination for the Smart Electricity Grid
18 Jeroen Janssens (UvT), Outlier Selection and One-Class Classification
19 Renze Steenhuisen (TUD), Coordinated Multi-Agent Planning and Scheduling
20 Katja Hofmann (UvA), Fast and Reliable Online Learning to Rank for Information Retrieval
21 Sander Wubben (UvT), Text-to-text generation by monolingual machine translation
22 Tom Claassen (RUN), Causal Discovery and Logic
23 Patricio de Alencar Silva (UvT), Value Activity Monitoring
24 Haitham Bou Ammar (UM), Automated Transfer in Reinforcement Learning
25 Agnieszka Anna Latoszek-Berendsen (UM), Intention-based Decision Support. A new way of representing and implementing clinical guidelines in a Decision Support System
26 Alireza Zarghami (UT), Architectural Support for Dynamic Homecare Service Provisioning
27 Mohammad Huq (UT), Inference-based Framework Managing Data Provenance
28 Frans van der Sluis (UT), When Complexity becomes Interesting: An Inquiry into the Information eXperience
29 Iwan de Kok (UT), Listening Heads
30 Joyce Nakatumba(TUE),Resource-Aware Business Process Management:Analysis and Support
31 Dinh Khoa Nguyen (UvT), Blueprint Model and Language for Engineering Cloud Applications
32 Kamakshi Rajagopal (OUN), Networking For Learning; The role of Networking in a Lifelong Learner's Professional Development
33 Qi Gao (TUD), User Modeling and Personalization in the Microblogging Sphere
34 Kien Tjin-Kam-Jet (UT), Distributed Deep Web Search
35 Abdallah El Ali (UvA), Minimal Mobile Human Computer Interaction
36 Than Lam Hoang (TUE), Pattern Mining in Data Streams
37 Dirk Börner (OUN), Ambient Learning Displays
38 Eelco den Heijer (VU), Autonomous Evolutionary Art
39 Joop de Jong (TUD), A Method for Enterprise Ontology based Design of Enterprise Information Systems
40 Pim Nijssen (UM), Monte-Carlo Tree Search for Multi-Player Games
41 Jochem Liem (UVA), Supporting the Conceptual Modelling of Dynamic Systems: A Knowledge Engineering Perspective on Qualitative Reasoning

- 42 Léon Planken (TUD), Algorithms for Simple Temporal Reasoning
43 Marc Bron (UVA), Exploration and Contextualization through Interaction and Concepts
-
- 2014 01 Nicola Barile (UU), Studies in Learning Monotone Models from Data
02 Fiona Tulyano (RUN), Combining System Dynamics with a Domain Modeling Method
03 Sergio Raul Duarte Torres (UT), Information Retrieval for Children: Search Behavior and Solutions
04 Hanna Jochmann-Mannak (UT), Websites for children: search strategies and interface design - Three studies on children's search performance and evaluation
05 Jurriaan van Reijssen (UU), Knowledge Perspectives on Advancing Dynamic Capability
06 Damian Tamburri (VU), Supporting Networked Software Development
07 Arya Adriansyah (TUE), Aligning Observed and Modeled Behavior
08 Samur Araujo (TUD), Data Integration over Distributed and Heterogeneous Data Endpoints
09 Philip Jackson (UvT), Toward Human-Level Artificial Intelligence: Representation and Computation of Meaning in Natural Language
10 Ivan Salvador Razo Zapata (VU), Service Value Networks
11 Janneke van der Zwaan (TUD), An Empathic Virtual Buddy for Social Support
12 Willem van Willigen (VU), Look Ma, No Hands: Aspects of Autonomous Vehicle Control
13 Arlette van Wissen (VU), Agent-Based Support for Behavior Change: Models and Applications in Health and Safety Domains
14 Yangyang Shi (TUD), Language Models With Meta-information
15 Natalya Mogles (VU), Agent-Based Analysis and Support of Human Functioning in Complex Socio-Technical Systems: Applications in Safety and Healthcare
16 Krystyna Milian (VU), Supporting trial recruitment and design by automatically interpreting eligibility criteria
17 Kathrin Dentler (VU), Computing healthcare quality indicators automatically: Secondary Use of Patient Data and Semantic Interoperability
18 Mattijs Ghijsen (UVA), Methods and Models for the Design and Study of Dynamic Agent Organizations
19 Vinicius Ramos (TUE), Adaptive Hypermedia Courses: Qualitative and Quantitative Evaluation and Tool Support
20 Mena Habib (UT), Named Entity Extraction and Disambiguation for Informal Text: The Missing Link
21 Cassidy Clark (TUD), Negotiation and Monitoring in Open Environments
22 Marieke Peeters (UU), Personalized Educational Games - Developing agent-supported scenario-based training
23 Eleftherios Sidirourgos (UvA/CWI), Space Efficient Indexes for the Big Data Era
24 Davide Ceolin (VU), Trusting Semi-structured Web Data
25 Martijn Lappenschaar (RUN), New network models for the analysis of disease interaction
26 Tim Baarslag (TUD), What to Bid and When to Stop
27 Rui Jorge Almeida (EUR), Conditional Density Models Integrating Fuzzy and Probabilistic Representations of Uncertainty
28 Anna Chmielowiec (VU), Decentralized k-Clique Matching
29 Jaap Kabbedijk (UU), Variability in Multi-Tenant Enterprise Software
30 Peter de Cock (UvT), Anticipating Criminal Behaviour
31 Leo van Moergestel (UU), Agent Technology in Agile Multiparallel Manufacturing and Product Support
32 Naser Ayat (UvA), On Entity Resolution in Probabilistic Data
33 Tesfa Tegegne (RUN), Service Discovery in eHealth
34 Christina Manteli (VU), The Effect of Governance in Global Software Development: Analyzing Transactive Memory Systems.
35 Joost van Ooijen (UU), Cognitive Agents in Virtual Worlds: A Middleware Design Approach
36 Joos Buijs (TUE), Flexible Evolutionary Algorithms for Mining Structured Process Models
37 Maral Dadvar (UT), Experts and Machines United Against Cyberbullying
38 Danny Plass-Oude Bos (UT), Making brain-computer interfaces better: improving usability through post-processing.
39 Jasmina Maric (UvT), Web Communities, Immigration, and Social Capital
40 Walter Omona (RUN), A Framework for Knowledge Manag. Using ICT in Higher Education
41 Frederic Hogenboom (EUR), Automated Detection of Financial Events in News Text
42 Carsten Eijckhof (CWI/TUD), Contextual Multidimensional Relevance Models
43 Kevin Vlaanderen (UU), Supporting Process Improvement using Method Increments
44 Paulien Meesters (UvT), Intelligent Blauw. Met als ondertitel: Intelligence-gestuurde politiezorg in gebiedsgebonden eenheden.
45 Birgit Schmitz (OUN), Mobile Games for Learning: A Pattern-Based Approach
46 Ke Tao (TUD), Social Web Data Analytics: Relevance, Redundancy, Diversity
47 Shangsong Liang (UVA), Fusion and Diversification in Information Retrieval
-
- 2015 01 Niels Netten (UvA), Machine Learning for Relevance of Information in Crisis Response
02 Faiza Bukhsh (UvT), Smart auditing: Innovative Compliance Checking in Customs Controls

-
- 03 Twan van Laarhoven (RUN), Machine learning for network data
 - 04 Howard Spoelstra (OUN), Collaborations in Open Learning Environments
 - 05 Christoph Bösch (UT), Cryptographically Enforced Search Pattern Hiding
 - 06 Farideh Heidari (TUD), Business Process Quality Computation - Computing Non-Functional Requirements to Improve Business Processes
 - 07 Maria-Hendrike Peetz (UvA), Time-Aware Online Reputation Analysis
 - 08 Jie Jiang (TUD), Organizational Compliance: An agent-based model for designing and evaluating organizational interactions
 - 09 Randy Klaassen (UT), HCI Perspectives on Behavior Change Support Systems
 - 10 Henry Hermans (OUN), OpenU: design of an integrated system to support lifelong learning
 - 11 Yongming Luo (TUE), Designing algorithms for big graph datasets: A study of computing bisimulation and joins
 - 12 Julie M. Birkholz (VU), Modi Operandi of Social Network Dynamics: The Effect of Context on Scientific Collaboration Networks
 - 13 Giuseppe Procaccianti (VU), Energy-Efficient Software
 - 14 Bart van Straalen (UT), A cognitive approach to modeling bad news conversations
 - 15 Klaas Andries de Graaf (VU), Ontology-based Software Architecture Documentation
 - 16 Changyun Wei (UT), Cognitive Coordination for Cooperative Multi-Robot Teamwork
 - 17 André van Cleeff (UT), Physical and Digital Security Mechanisms: Properties, Combinations and Trade-offs
 - 18 Holger Pirk (CWI), Waste Not, Want Not! Managing Relational Data in Asymmetric Memories
 - 19 Bernardo Tabuenca (OUN), Ubiquitous Technology for Lifelong Learners
 - 20 Lois Vanhée (UU), Using Culture and Values to Support Flexible Coordination
 - 21 Sibren Fetter (OUN), Using Peer-Support to Expand and Stabilize Online Learning
 - 22 Zheming Zhu (UT), Co-occurrence Rate Networks
 - 23 Luit Gazendam (VU), Cataloguer Support in Cultural Heritage
 - 24 Richard Berendsen (UVA), Finding People, Papers, and Posts: Vertical Search Algorithms and Evaluation
 - 25 Steven Woudenberg (UU), Bayesian Tools for Early Disease Detection
 - 26 Alexander Hogenboom (EUR), Sentiment Analysis of Text Guided by Semantics and Structure
 - 27 Sándor Héman (CWI), Updating compressed column stores
 - 28 Janet Bagorogoza (TiU), Knowledge Management and High Performance; The Uganda Financial Institutions Model for HPO
 - 29 Hendrik Baier (UM), Monte-Carlo Tree Search Enhancements for One-Player and Two-Player Domains
 - 30 Kiavash Bahreini (OU), Real-time Multimodal Emotion Recognition in E-Learning
 - 31 Yakup Koç (TUD), On the robustness of Power Grids
 - 32 Jerome Gard (UL), Corporate Venture Management in SMEs
 - 33 Frederik Schadd (TUD), Ontology Mapping with Auxiliary Resources
 - 34 Victor de Graaf (UT), Gesocial Recommender Systems
 - 35 Jungxao Xu (TUD), Affective Body Language of Humanoid Robots: Perception and Effects in Human Robot Interaction
-
- 2016 01 Syed Saiden Abbas (RUN), Recognition of Shapes by Humans and Machines
 - 02 Michiel Christiaan Meulendijk (UU), Optimizing medication reviews through decision support: prescribing a better pill to swallow
 - 03 Maya Sappelli (RUN), Knowledge Work in Context: User Centered Knowledge Worker Support
 - 04 Laurens Rietveld (VU), Publishing and Consuming Linked Data
 - 05 Evgeny Sherkhonov (UVA), Expanded Acyclic Queries: Containment and an Application in Explaining Missing Answers
 - 06 Michel Wilson (TUD), Robust scheduling in an uncertain environment
 - 07 Jeroen de Man (VU), Measuring and modeling negative emotions for virtual training
 - 08 Matje van de Camp (TiU), A Link to the Past: Constructing Historical Social Networks from Unstructured Data
 - 09 Archana Nottamkandath (VU), Trusting Crowdsourced Information on Cultural Artefacts
 - 10 George Karafotias (VUA), Parameter Control for Evolutionary Algorithms
 - 11 Anne Schuth (UVA), Search Engines that Learn from Their Users
 - 12 Max Knobbout (UU), Logics for Modelling and Verifying Normative Multi-Agent Systems
 - 13 Nana Baah Gyan (VU), The Web, Speech Technologies and Rural Development in West Africa - An ICT4D Approach
 - 14 Ravi Khadka (UU), Revisiting Legacy Software System Modernization
 - 15 Steffen Michels (RUN), Hybrid Probabilistic Logics - Theoretical Aspects, Algorithms and Experiments
 - 16 Guangliang Li (UVA), Socially Intelligent Autonomous Agents that Learn from Human Reward
 - 17 Berend Weel (VU), Towards Embodied Evolution of Robot Organisms
 - 18 Albert Meroño Peñuela (VU), Refining Statistical Data on the Web
 - 19 Julia Efremova (Tu/e), Mining Social Structures from Genealogical Data
 - 20 Daan Odijk (UVA), Context & Semantics in News & Web Search

- 21 Alejandro Moreno Céleri (UT), From Traditional to Interactive Playspaces: Automatic Analysis of Player Behavior in the Interactive Tag Playground
 - 22 Grace Lewis (VU), Software Architecture Strategies for Cyber-Foraging Systems
 - 23 Fei Cai (UVA), Query Auto Completion in Information Retrieval
 - 24 Brend Wanders (UT), Repurposing and Probabilistic Integration of Data; An Iterative and data model independent approach
 - 25 Julia Kiseleva (TU/e), Using Contextual Information to Understand Searching and Browsing Behavior
 - 26 Dilhan Thilakarathne (VU), In or Out of Control: Exploring Computational Models to Study the Role of Human Awareness and Control in Behavioural Choices, with Applications in Aviation and Energy Management Domains
 - 27 Wen Li (TUD), Understanding Geo-spatial Information on Social Media
 - 28 Mingxin Zhang (TUD), Large-scale Agent-based Social Simulation - A study on epidemic prediction and control
 - 29 Nicolas Höning (TUD), Peak reduction in decentralised electricity systems - Markets and prices for flexible planning
 - 30 Ruud Mattheij (UvT), The Eyes Have It
 - 31 Mohammad Khelghati (UT), Deep web content monitoring
 - 32 Eelco Vriezekolk (UT), Assessing Telecommunication Service Availability Risks for Crisis Organisations
 - 33 Peter Bloem (UVA), Single Sample Statistics, exercises in learning from just one example
 - 34 Dennis Schunselaar (TUE), Configurable Process Trees: Elicitation, Analysis, and Enactment
 - 35 Zhaochun Ren (UVA), Monitoring Social Media: Summarization, Classification and Recommendation
 - 36 Daphne Karreman (UT), Beyond R2D2: The design of nonverbal interaction behavior optimized for robot-specific morphologies
 - 37 Giovanni Sileno (UvA), Aligning Law and Action - a conceptual and computational inquiry
 - 38 Andrea Minuto (UT), Materials that Matter - Smart Materials meet Art & Interaction Design
 - 39 Merijn Bruijnes (UT), Believable Suspect Agents; Response and Interpersonal Style Selection for an Artificial Suspect
 - 40 Christian Detweiler (TUD), Accounting for Values in Design
 - 41 Thomas King (TUD), Governing Governance: A Formal Framework for Analysing Institutional Design and Enactment Governance
 - 42 Spyros Martzoukos (UVA), Combinatorial and Compositional Aspects of Bilingual Aligned Corpora
 - 43 Saskia Koldijk (RUN), Context-Aware Support for Stress Self-Management: From Theory to Practice
 - 44 Thibault Sellam (UVA), Automatic Assistants for Database Exploration
 - 45 Bram van de Laar (UT), Experiencing Brain-Computer Interface Control
 - 46 Jorge Gallego Perez (UT), Robots to Make you Happy
 - 47 Christina Weber (UL), Real-time foresight - Preparedness for dynamic innovation networks
 - 48 Tanja Buttler (TUD), Collecting Lessons Learned
 - 49 Gleb Polevoy (TUD), Participation and Interaction in Projects. A Game-Theoretic Analysis
 - 50 Yan Wang (UvT), The Bridge of Dreams: Towards a Method for Operational Performance Alignment in IT-enabled Service Supply Chains
-
- 2017 01 Jan-Jaap Oerlemans (UL), Investigating Cybercrime
 - 02 Sjoerd Timmer (UU), Designing and Understanding Forensic Bayesian Networks using Argumentation
 - 03 Daniël Harold Telgen (UU), Grid Manufacturing; A Cyber-Physical Approach with Autonomous Products and Reconfigurable Manufacturing Machines
 - 04 Mrunal Gawade (CWI), Multi-core Parallelism in a Column-store
 - 05 Mahdiah Shadi (UVA), Collaboration Behavior
 - 06 Damir Vadic (EUR), Intelligent Information Systems for Web Product Search
 - 07 Roel Bertens (UU), Insight in Information: from Abstract to Anomaly
 - 08 Rob Konijn (VU) , Detecting Interesting Differences: Data Mining in Health Insurance Data using Outlier Detection and Subgroup Discovery
 - 09 Dong Nguyen (UT), Text as Social and Cultural Data: A Computational Perspective on Variation in Text
 - 10 Robby van Delden (UT), (Steering) Interactive Play Behavior
 - 11 Florian Kunneman (RUN), Modelling patterns of time and emotion in Twitter
 - 12 Sander Leemans (TUE), Robust Process Mining with Guarantees
 - 13 Gijs Huisman (UT), Social Touch Technology - Extending the reach of social touch through haptic technology
 - 14 Shoshannah Tekofsky (UvT), You Are Who You Play You Are: Modelling Player Traits from Video Game Behavior
 - 15 Peter Berck (RUN), Memory-Based Text Correction
 - 16 Aleksandr Chuklin (UVA), Understanding and Modeling Users of Modern Search Engines
 - 17 Daniel Dimov (UL), Crowdsourced Online Dispute Resolution

-
- 18 Ridho Reinanda (UVA), Entity Associations for Search
 - 19 Jeroen Vuurens (UT), Proximity of Terms, Texts and Semantic Vectors in Information Retrieval
 - 20 Mohammadbashir Sedighi (TUD), Fostering Engagement in Knowledge Sharing: The Role of Perceived Benefits, Costs and Visibility
 - 21 Jeroen Linssen (UT), Meta Matters in Interactive Storytelling and Serious Gaming (A Play on Worlds)
 - 22 Sara Magliacane (VU), Logics for causal inference under uncertainty
 - 23 David Graus (UVA), Entities of Interest — Discovery in Digital Traces
 - 24 Chang Wang (TUD), Use of Affordances for Efficient Robot Learning
 - 25 Veruska Zamborlini (VU), Knowledge Representation for Clinical Guidelines, with applications to Multimorbidity Analysis and Literature Search
 - 26 Merel Jung (UT), Socially intelligent robots that understand and respond to human touch
 - 27 Michiel Joesse (UT), Investigating Positioning and Gaze Behaviors of Social Robots: People's Preferences, Perceptions and Behaviors
 - 28 John Klein (VU), Architecture Practices for Complex Contexts
 - 29 Adel Alhuraibi (UvT), From IT-BusinessStrategic Alignment to Performance: A Moderated Mediation Model of Social Innovation, and Enterprise Governance of IT"
 - 30 Wilma Latuny (UvT), The Power of Facial Expressions
 - 31 Ben Ruijl (UL), Advances in computational methods for QFT calculations
 - 32 Thaeer Samar (RUN), Access to and Retrieval of Content in Web Archives
 - 33 Brigit van Loggem (OU), Towards a Design Rationale for Software Documentation: A Model of Computer-Mediated Activity
 - 34 Maren Scheffel (OU), The Evaluation Framework for Learning Analytics
 - 35 Martine de Vos (VU), Interpreting natural science spreadsheets
 - 36 Yuanhao Guo (UL), Shape Analysis for Phenotype Characterisation from High-throughput Imaging
 - 37 Alejandro Montes Garcia (TUE), WiBAF: A Within Browser Adaptation Framework that Enables Control over Privacy
 - 38 Alex Kayal (TUD), Normative Social Applications
 - 39 Sara Ahmadi (RUN), Exploiting properties of the human auditory system and compressive sensing methods to increase noise robustness in ASR
 - 40 Altaf Hussain Abro (VUA), Steer your Mind: Computational Exploration of Human Control in Relation to Emotions, Desires and Social Support For applications in human-aware support systems
 - 41 Adnan Manzoor (VUA), Minding a Healthy Lifestyle: An Exploration of Mental Processes and a Smart Environment to Provide Support for a Healthy Lifestyle
 - 42 Elena Sokolova (RUN), Causal discovery from mixed and missing data with applications on ADHD datasets
 - 43 Maaike de Boer (RUN), Semantic Mapping in Video Retrieval
 - 44 Garm Lucassen (UU), Understanding User Stories - Computational Linguistics in Agile Requirements Engineering
 - 45 Bas Testerink (UU), Decentralized Runtime Norm Enforcement
 - 46 Jan Schneider (OU), Sensor-based Learning Support
 - 47 Jie Yang (TUD), Crowd Knowledge Creation Acceleration
 - 48 Angel Suarez (OU), Collaborative inquiry-based learning
-
- 2018 01 Han van der Aa (VUA), Comparing and Aligning Process Representations
 - 02 Felix Mannhardt (TUE), Multi-perspective Process Mining
 - 03 Steven Bosems (UT), Causal Models For Well-Being: Knowledge Modeling, Model-Driven Development of Context-Aware Applications, and Behavior Prediction
 - 04 Jordan Janeiro (TUD), Flexible Coordination Support for Diagnosis Teams in Data-Centric Engineering Tasks
 - 05 Hugo Hurdeman(UVA),Supporting the Complex Dynamics of the Information Seeking Process
 - 06 Dan Ionita(UT),Model-Driven Inform. Security Risk Assessment of SocioTechnical Systems
 - 07 Jieting Luo (UU), A formal account of opportunism in multi-agent systems
 - 08 Rick Smeters (RUN), Advances in Model Learning for Software Systems
 - 09 Xu Xie (TUD), Data Assimilation in Discrete Event Simulations
 - 10 Julienka Mollee (VUA), Moving forward: supporting physical activity behavior change through intelligent technology
 - 11 Mahdi Sargolzaei (UVA), Enabling Framework for Service-oriented Collaborative Networks
 - 12 Xixi Lu (TUE), Using behavioral context in process mining
 - 13 Seyed Amin Tabatabaei (VUA), Computing a Sustainable Future
 - 14 Bart Joosten (UvT), Detecting Social Signals with Spatiotemporal Gabor Filters
 - 15 Naser Davarzani (UM), Biomarker discovery in heart failure
 - 16 Jaebok Kim (UT), Automatic recognition of engagement and emotion in a group of children
 - 17 Jianpeng Zhang (TUE), On Graph Sample Clustering
 - 18 Henriette Nakad (UL), De Notaris en Private Rechtspraak
 - 19 Minh Duc Pham (VUA), Emergent relational schemas for RDF

- 20 Manxia Liu (RUN), Time and Bayesian Networks
 - 21 Aad Sloodmaker (OUN), EMERGO: a generic platform for authoring and playing scenario-based serious games
 - 22 Eric Fernandes de Mello Araujo (VUA), Contagious: Modeling the Spread of Behaviours, Perceptions and Emotions in Social Networks
 - 23 Kim Schouten (EUR), Semantics-driven Aspect-Based Sentiment Analysis
 - 24 Jered Vroon (UT), Responsive Social Positioning Behaviour for Semi-Autonomous Telepresence Robots
 - 25 Riste Gligorov (VUA), Serious Games in Audio-Visual Collections
 - 26 Roelof Anne Jelle de Vries (UT), Theory-Based and Tailor-Made: Motivational Messages for Behavior Change Technology
 - 27 Maikel Leemans (TUE), Hierarchical Process Mining for Scalable Software Analysis
 - 28 Christian Willemse (UT), Social Touch Technologies: How they feel and how they make you feel
 - 29 Yu Gu (UVT), Emotion Recognition from Mandarin Speech
 - 30 Wouter Beek, The "K" in "semantic web" stands for "knowledge": scaling semantics to the web
-
- 2019 01 Rob van Eijk (UL), Web privacy measurement in real-time bidding systems. A graph-based approach to RTB system classification
 - 02 Emmanuelle Beauxis Aussalet (CWI, UU), Statistics and Visualizations for Assessing Class Size Uncertainty
 - 03 Eduardo Gonzalez Lopez de Murillas (TUE), Process Mining on Databases: Extracting Event Data from Real Life Data Sources
 - 04 Ridho Rahmadi (RUN), Finding stable causal structures from clinical data
 - 05 Sebastiaan van Zelst (TUE), Process Mining with Streaming Data
 - 06 Chris Dijkshoorn (VU), Nichesourcing for Improving Access to Linked Cultural Heritage Datasets
 - 07 Soude Fazeli (TUD), Recommender Systems in Social Learning Platforms
 - 08 Frits de Nijs (TUD), Resource-constrained Multi-agent Markov Decision Processes
 - 09 Fahimeh Alizadeh Moghaddam (UVA), Self-adaptation for energy efficiency in software systems
 - 10 Qing Chuan Ye (EUR), Multi-objective Optimization Methods for Allocation and Prediction
 - 11 Yue Zhao (TUD), Learning Analytics Technology to Understand Learner Behavioral Engagement in MOOCs
 - 12 Jacqueline Heijerman (VU), Better Together
 - 13 Guanliang Chen (TUD), MOOC Analytics: Learner Modeling and Content Generation
 - 14 Daniel Davis (TUD), Large-Scale Learning Analytics: Modeling Learner Behavior & Improving Learning Outcomes in Massive Open Online Courses
 - 15 Erwin Walraven (TUD), Planning under Uncertainty in Constrained and Partially Observable Environments
 - 16 Guangming Li (TUE), Process Mining based on Object-Centric Behavioral Constraint (OCBC) Models
 - 17 Ali Hurriyetoglu (RUN), Extracting actionable information from microtexts
 - 18 Gerard Wagenaar (UU), Artefacts in Agile Team Communication
 - 19 Vincent Koeman (TUD), Tools for Developing Cognitive Agents
 - 20 Chide Groenouwe (UU), Fostering technically augmented human collective intelligence
 - 21 Cong Liu (TUE), Software Data Analytics: Architectural Model Discovery and Design Pattern Detection
 - 22 Martin van den Berg (VU), Improving IT Decisions with Enterprise Architecture
 - 23 Qin Liu (TUD), Intelligent Control Systems: Learning, Interpreting, Verification
 - 24 Anca Dumitrache (VU), Truth in Disagreement - Crowdsourcing Labeled Data for Natural Language Processing
 - 25 Emiel van Miltenburg (VU), Pragmatic factors in (automatic) image description
 - 26 Prince Singh (UT), An Integration Platform for Synchronodal Transport
 - 27 Alessandra Antonaci (OUN), The Gamification Design Process applied to (Massive) Open Online Courses
 - 28 Esther Kuindersma (UL), Cleared for take-off: Game-based learning to prepare airline pilots for critical situations
 - 29 Daniel Formolo (VU), Using virtual agents for simulation and training of social skills in safety-critical circumstances
 - 30 Vahid Yazdanpanah (UT), Multiagent Industrial Symbiosis Systems
 - 31 Milan Jelisavcic (VU), Alive and Kicking: Baby Steps in Robotics
 - 32 Chiara Sironi (UM), Monte-Carlo Tree Search for Artificial General Intelligence in Games
 - 33 Anil Yaman (TUE), Evolution of Biologically Inspired Learning in Artificial Neural Networks
 - 34 Negar Ahmadi (TUE), EEG Microstate and Functional Brain Network Features for Classification of Epilepsy and PNES
 - 35 Lisa Facey-Shaw (OUN), Gamification with digital badges in learning programming
 - 36 Kevin Ackermans (OUN), Designing Video-Enhanced Rubrics to Master Complex Skills
 - 37 Jian Fang (TUD), Database Acceleration on FPGAs
 - 38 Akos Kadar (OUN), Learning visually grounded and multilingual representations
-

-
- 2020 01 Armon Toubman (UL), Calculated Moves: Generating Air Combat Behaviour
 02 Marcos de Paula Bueno (UL), Unraveling Temporal Processes using Probabilistic Graphical Models
 03 Mostafa Deghani (UvA), Learning with Imperfect Supervision for Language Understanding
 04 Maarten van Gompel (RUN), Context as Linguistic Bridges
 05 Yulong Pei (TUE), On local and global structure mining
 06 Preethu Rose Anish (UT), Stimulation Architectural Thinking during Requirements Elicitation - An Approach and Tool Support
 07 Wim van der Vegt (OUN), Towards a software architecture for reusable game components
 08 Ali Mirsoleimani (UL), Structured Parallel Programming for Monte Carlo Tree Search
 09 Myriam Traub (UU), Measuring Tool Bias and Improving Data Quality for Digital Humanities Research
 10 Alifah Syamsiyah (TUE), In-database Preprocessing for Process Mining
 11 Sepideh Mesbah (TUD), Semantic-Enhanced Training Data Augmentation Methods for Long-Tail Entity Recognition Models
 12 Ward van Breda (VU), Predictive Modeling in E-Mental Health: Exploring Applicability in Personalised Depression Treatment
 13 Marco Virgolin (CWI), Design and Application of Gene-pool Optimal Mixing Evolutionary Algorithms for Genetic Programming
 14 Mark Raasveldt (CWI/UL), Integrating Analytics with Relational Databases
 15 Konstantinos Georgiadis (OUN), Smart CAT: Machine Learning for Configurable Assessments in Serious Games
 16 Ilona Wilmont (RUN), Cognitive Aspects of Conceptual Modelling
 17 Daniele Di Mitri (OUN), The Multimodal Tutor: Adaptive Feedback from Multimodal Experiences
 18 Georgios Methenitis (TUD), Agent Interactions & Mechanisms in Markets with Uncertainties: Electricity Markets in Renewable Energy Systems
 19 Guido van Capelleveen (UT), Industrial Symbiosis Recommender Systems
 20 Albert Hankel (VU), Embedding Green ICT Maturity in Organisations
 21 Karine da Silva Miras de Araujo (VU), Where is the robot?: Life as it could be
 22 Maryam Masoud Khamis (RUN), Understanding complex systems implementation through a modeling approach: the case of e-government in Zanzibar
 23 Rianne Conijn (UT), The Keys to Writing: A writing analytics approach to studying writing processes using keystroke logging
 24 Lenin da Nobrega Medeiros (VUA/RUN), How are you feeling, human? Towards emotionally supportive chatbots
 25 Xin Du (TUE), The Uncertainty in Exceptional Model Mining
 26 Krzysztof Leszek Sadowski (UU), GAMBIT: Genetic Algorithm for Model-Based mixed-Integer Optimization
 27 Ekaterina Muravyeva (TUD), Personal data and informed consent in an educational context
 28 Bibeg Limbu (TUD), Multimodal interaction for deliberate practice: Training complex skills with augmented reality
 29 Ioan Gabriel Bucur (RUN), Being Bayesian about Causal Inference
 30 Bob Zadok Blok (UL), Creatief, Creatieve, Creatiefst
 31 Gongjin Lan (VU), Learning better - From Baby to Better
 32 Jason Rhuggenaath (TUE), Revenue management in online markets: pricing and online advertising
 33 Rick Gilsing (TUE), Supporting service-dominant business model evaluation in the context of business model innovation
 34 Anna Bon (MU), Intervention or Collaboration? Redesigning Information and Communication Technologies for Development
 35 Siamak Farshidi (UU), Multi-Criteria Decision-Making in Software Production
-
- 2021 01 Francisco Xavier Dos Santos Fonseca (TUD), Location-based Games for Social Interaction in Public Space
 02 Rijk Mercuur (TUD), Simulating Human Routines: Integrating Social Practice Theory in Agent-Based Models
 03 Seyyed Hadi Hashemi (UVA), Modeling Users Interacting with Smart Devices
 04 Ioana Jivet (OU), The Dashboard That Loved Me: Designing adaptive learning analytics for self-regulated learning
 05 Davide Dell'Anna (UU), Data-Driven Supervision of Autonomous Systems
 06 Daniel Davison (UT), "Hey robot, what do you think?" How children learn with a social robot
 07 Arnel Lefebvre (UU), Research data management for open science
 08 Nardie Fanchamps (OU), The Influence of Sense-Reason-Act Programming on Computational Thinking
 09 Cristina Zaga (UT), The Design of Robotings. Non-Anthropomorphic and Non-Verbal Robots to Promote Children's Collaboration Through Play
 10 Quinten Meertens (UvA), Misclassification Bias in Statistical Learning
 11 Anne van Rossum (UL), Nonparametric Bayesian Methods in Robotic Vision

- 12 Lei Pi (UL), External Knowledge Absorption in Chinese SMEs
 - 13 Bob R. Schadenberg (UT), Robots for Autistic Children: Understanding and Facilitating Predictability for Engagement in Learning
 - 14 Negin Samaeemofrad (UL), Business Incubators: The Impact of Their Support
 - 15 Onat Ege Adali (TU/e), Transformation of Value Propositions into Resource Re-Configurations through the Business Services Paradigm
 - 16 Esam A. H. Ghaleb (UM), BIMODAL EMOTION RECOGNITION FROM AUDIO-VISUAL CUES
 - 17 Dario Dotti (UM), Human Behavior Understanding from motion and bodily cues using deep neural networks
 - 18 Remi Wieten (UU), Bridging the Gap Between Informal Sense-Making Tools and Formal Systems - Facilitating the Construction of Bayesian Networks and Argumentation Frameworks
 - 19 Roberto Verdecchia (VU), Architectural Technical Debt: Identification and Management
 - 20 Masoud Mansoury (TU/e), Understanding and Mitigating Multi-Sided Exposure Bias in Recommender Systems
 - 21 Pedro Thiago Timbó Holanda (CWI), Progressive Indexes
 - 22 Sihang Qiu (TUD), Conversational Crowdsourcing
 - 23 Hugo Manuel Proença (LIACS), Robust rules for prediction and description
 - 24 Kaijie Zhu (TUE), On Efficient Temporal Subgraph Query Processing
 - 25 Eoin Martino Grua (VUA), The Future of E-Health is Mobile: Combining AI and Self-Adaptation to Create Adaptive E-Health Mobile Applications
 - 26 Benno Kruit (CWI & VUA), Reading the Grid: Extending Knowledge Bases from Human-readable Tables
 - 27 Jelte van Waterschoot (UT), Personalized and Personal Conversations: Designing Agents Who Want to Connect With You
 - 28 Christoph Selig(UL),Understanding the Heterogeneity of Corp. Entrepreneurship Programs
-
- 2022 1 Judith van Stegeren (UT), Flavor text generation for role-playing video games
 - 2 Paulo da Costa (TU/e), Data-driven Prognostics and Logistics Optimisation: A Deep Learning Journey
 - 3 Ali el Hassouni (VUA), A Model A Day Keeps The Doctor Away: Reinforcement Learning For Personalized Healthcare
 - 4 Unal Aksu (UU), A Cross-Organizational Process Mining Framework
 - 5 Shiwei Liu (TU/e), Sparse Neural Network Training with In-Time Over-Parameterization
 - 6 Reza Refaei Afshar (TU/e), Machine Learning for Ad Publishers in Real Time Bidding
 - 7 Sambit Praharaaj (OU), Measuring the Unmeasurable? Towards Automatic Co-located Collaboration Analytics
 - 8 Maikel L. van Eck (TU/e), Process Mining for Smart Product Design
 - 9 Oana Andreea Inel (VUA), Understanding Events: A Diversity-driven Human-Machine Approach
 - 10 Felipe Moraes Gomes (TUD), Examining the Effectiveness of Collaborative Search Engines
 - 11 Mirjam de Haas (UT), Staying engaged in child-robot interaction, a quantitative approach to studying preschoolers engagement with robots and tasks during second-language tutoring
 - 12 Guanyi Chen (UU), Computational Generation of Chinese Noun Phrases
 - 13 Xander Wilcke (VUA), Machine Learning on Multimodal Knowledge Graphs: Opportunities, Challenges, and Methods for Learning on Real-World Heterogeneous and Spatially-Oriented Knowledge
 - 14 Michiel Overeem (UU), Evolution of Low-Code Platforms
 - 15 Jelmer Jan Koorn (UU), Work in Process: Unearthing Meaning using Process Mining
 - 16 Pieter Gijbbers (TU/e), Systems for AutoML Research
 - 17 Laura van der Lubbe(VUA),Empowering vulnerable people w. serious games and gamification
 - 18 Paris Mavromoustakos Blom (TiU), Player Affect Modelling and Video Game Personalisation
 - 19 Bilge Yigit Ozkan (UU), Cybersecurity Maturity Assessment and Standardisation
 - 20 Fakhra Jabeen (VUA), Dark Side of the Digital Media - Computational Analysis of Negative Human Behaviors on Social Media
 - 21 Seethu Mariyam Christopher (UM), Intelligent Toys for Physical and Cognitive Assessments
 - 22 Alexandra Sierra Rativa (TiU), Virtual Character Design and its potential to foster Empathy, Immersion, and Collaboration Skills in Video Games and Virtual Reality Simulations
 - 23 Ilir Kola (TUD), Enabling Social Situation Awareness in Support Agents
 - 24 Samaneh Heidari (UU), Agents with Social Norms and Values - A framework for agent based social simulations with social norms and personal values
 - 25 Anna L.D. Latour (LU), Optimal decision-making under constraints and uncertainty
 - 26 Anne Dirkson (LU), Knowledge Discovery from Patient Forums: Gaining novel medical insights from patient experiences
 - 27 Christos Athanasiadis(UM),Emotion-aware cross-modal domain adaptation in video sequences
 - 28 Onuralp Ulusoy (UU), Privacy in Collaborative Systems
 - 29 Jan Kolkmeier (UT), From Head Transform to Mind Transplant: Social Interactions in Mixed Reality
-