

PACCART: Reinforcing Trust in Multiuser Privacy Agreement Systems (Full)

Daan Di Scala¹[0000–0003–1548–6675] and Pinar Yolum²[0000–0001–7848–1834]

¹ Utrecht University, Utrecht, The Netherlands, daandiscala@hotmail.com

² Utrecht University, Utrecht, The Netherlands, p.yolum@uu.nl

Abstract. Collaborative systems, such as Online Social Networks and the Internet of Things, enable users to share privacy sensitive content. Content in these systems is often co-owned by multiple users with different privacy expectations, leading to possible multiuser privacy conflicts. In order to resolve these conflicts, various agreement mechanisms have been designed and agents that could participate in such mechanisms have been proposed. However, research shows that users hesitate to use software tools for managing their privacy. To remedy this, we argue that users should be supported by trustworthy agents that adhere to the following criteria: (i) concealment of privacy preferences, such that only necessary information is shared with others, (ii) equity of treatment, such that different kinds of users are supported equally, (iii) collaboration of users, such that a group of users can support each other in agreement and (iv) explainability of actions, such that users know why certain information about them was shared to reach a decision. Accordingly, this paper proposes PACCART, an open-source agent that satisfies these criteria. Our experiments over simulations and user study indicate that PACCART increases user trust significantly.

Keywords: Multiuser Privacy · Trust · Equity.

1 Introduction

Privacy is the right of individuals to keep personal information to themselves [31]. While many systems are built with configurations to enable users to exercise this right, managing privacy is still a difficult problem. On one hand, collaborative systems, such as Online Social Networks and Internet of Things, contain a vast amount of content that pertain to a single individual, making it difficult, if not impossible, for individuals to attend to each piece of content separately [20]. Recent research on privacy agents shows promising results on how agents can help with privacy, such as on detecting privacy violations [14], recommending sharing behavior [11,26], and learning privacy preferences [16,30]. An important aspect to consider is **co-owned** content, such that the content does not belong to a single individual (e.g., medical information), but pertains to multiple people (e.g., a group photo or co-edited document [10]). These co-owners of the content

can and do have conflicting desires about the usage of the content, leading to what is termed as **multiuser privacy conflicts (MPCs)** [23,28].

Various decision-making techniques, such as auctions, negotiation, and argumentation have been employed to build systems to resolve MPCs. Simply put, each user that participates in these systems is represented by a privacy agent that knows its user’s privacy requirements. The agent participates in the decision-making system on behalf of its user. For auction-based systems, this means bidding on its user’s behalf or for argumentation-based systems, this would correspond generating arguments on behalf of its user. Through participation in this system, the agents decide if and how to share co-owned content by resolving conflicts. Experimental evaluations on these systems yield good performance results. However, it is also known that users have concerns when it comes to using software tools for managing various elements of their privacy [12,27].

Many existing studies of collaborative systems indicate the importance of *trust* in making systems usable by individuals [5,17]. We argue that to realize trust, the privacy agent of a user should satisfy the following properties:

Concealment: The privacy agent will know the privacy constraints of the user, either through elicitation or learning over time. When the agent is interacting with others to resolve conflicts, it should reveal as little as possible about these privacy constraints, since the privacy constraints themselves are private information. So, users would know that their privacy is safe with their agent [2,17].

Equity: Different users have different privacy stances in terms of their motivation and knowledge. While some users would fight not to share a piece of content, others will be indifferent. Contrary to some of the existing work in AI that favors users with certain properties [19,24], we do not want any user to be left behind. Ideally, the privacy agent should take the privacy stance of the user into account and be able to help different types of users as equally as possible; thereby creating equity [31,33].

Collaboration: It is possible that a number of agents that participate in the same conflict resolution have similar privacy concerns or complementary information to support a particular privacy decision [32]. Their agents should be able to collaborate in groups.

Explainability: It is well-studied that often users do not trust privacy tools because of misconceptions [27]. One solution for this is to make the tools explicit to users. But, more importantly, if the agent itself can provide explanations as to why it has taken certain actions, then its user can understand and even configure the agent better for future interactions [9,21].

Accordingly, this paper proposes a new Privacy Agent for Content Concealment in Argumentation to Reinforce Trust (PACCART). PACCART can conceal its user’s privacy requirements at different levels, while still resolving conflicts. By adapting to different privacy understandings of users, PACCART will provide equitable treatment. At the same time, PACCART will enable agents to work together towards a shared desired outcome. Finally, it will help its user understand the actions it is taking. To the best of our knowledge, this is the first

privacy agent that brings these desirable properties together. We made PACCART openly available.³

The rest of this paper is organized as follows: Section 2 explains the necessary background theory on argumentation-based agreement systems. Section 3 formalizes the PACCART model. Section 4 describes our realization of the model and our experimental results. Section 5 discusses the user study and its results. Finally, Section 6 systematically compares our approach with related work and gives pointers for future directions.

2 Background

We advocate that for an agent to exhibit these four criteria, it is useful to be able to express the relations between privacy preferences in a semantic manner. Thus, as an underlying agreement system, we opt for argumentation as opposed to other decision-making mechanisms such as auctions or negotiation. Below, we review how a privacy agent would use argumentation theory and how by using a dialogical argumentation system it can resolve privacy disputes.

2.1 Argumentation Theory

Our agent model makes use of argumentation theory for its reasoning. We follow the **structured argumentation** formalism of ASPIC+ [22]. An ASPIC+ **argumentation** or **dispute** $d = \langle P, R, B, C \rangle$ consists of **premises** $P = P_o \cup P_n$ (ordinary premises P_o and necessary premises P_n), **rules** $R = R_s \cup R_d$ (strict rules R_s and defeasible rules R_d), **biases** $B = B_p \cup B_r$ (premise biases B_p and rule biases B_r) and Contraries C .

A dispute is held between two opposing agents, **proponent** a_p and **opponent** a_o . Agents have access to their **knowledge base** KB , which contains premises, rules and contraries. With this content, agents can form **arguments**. In order to win the dispute, agents are able to **attack** each other’s arguments and can **support** (or **defend**) their own arguments with subarguments in order to try to win the dispute [7]. In some cases an agent is also able to **forfeit**, giving up on winning the dispute. Arguments can be attacked on their **weak points**, which is any subargument that is either a consequent of a defeasible rule or any ordinary premise. **Useful** arguments are arguments that, when added to the dispute, successfully attack any opponent’s current arguments. Acceptability conditions of winning or losing are dependent on the chosen **semantics**. Baroni et al. [3] offer an overview of different semantics and their meaning, including **grounded**, **preferred**, **complete** and **stable** semantics.

2.2 Dispute Protocol

In order for an argumentation agent to be able to hold a dispute with other agents about a **subject**, it follows a communication protocol. The protocol allows

³ <https://github.com/PACCART/PACCARTpaper>

agents to **extend** the dispute, meaning that they take turns adding arguments from their knowledge base to the dispute in order to either defend or attack the dispute subject.

Algorithm 1 Agent Dispute Extension Protocol

Require: Agents $A = \{a_p, a_o\}$, each with $KB = \langle P, R, C \rangle$

Ensure: Determine winner of dispute d

```

1:  $a \leftarrow a_p$ 
2: while  $d$  is not forfeited do
3:   if  $a$  can extend  $d$  then
4:      $a$  extends  $d$ 
5:     if  $a$  is  $a_p$  then
6:        $a \leftarrow a_o$ 
7:     else
8:        $a \leftarrow a_p$ 
9:   else
10:     $a$  forfeits  $d$ 

```

Argumentation systems like PriArg [13] utilize this kind of extension protocol, as denoted in Algorithm 1. According to the extension protocol, if an agent is able to extend the dispute, it does so. An agent extends the dispute by adding any sufficient argument from its knowledge base. Therefore, as soon as an agent is unable to extend the dispute any further, it forfeits the dispute.

The winner of a dispute is determined by evaluating the outcome according to grounded semantics. This way the **burden of proof** initially lies on the proponent of the dispute, after which agents take turns by extending the dispute until one of them wins. This is done because the agent that initializes the dispute has something to gain by defending the subject.

3 Model

The PACCART agent consists of a base component, which works similarly to agents in the PriArg system, as it communicates with other agents through a dialogical argumentation framework that follows the same Dispute Extension Protocol, as defined in Section 2.2. Following this, four components will be introduced on top of the workings of the base component.

3.1 Concealment Component

In the case of argumentation over privacy issues, the information to be concealed consists of all information that a user’s agent can hold in its knowledge base, including those that pertain to the user’s privacy preferences. We make a distinction between content that is revealed during a dispute and content that is not (yet) revealed, by keeping track of concealed rules R_c and premises P_c .

We make a distinction between content in Agent A 's knowledge base KB that is concealed and content that is not, by keeping track of different sets throughout the dispute. At the initialization stage of the dispute, agents have not yet shared any content with each other, which means that all content is still concealed ($R_c = R$ and $P_c = P$). While the dispute develops, each time an agent shares content with another agent to extend the dispute, that content is revealed and therefore removed from the set of concealed content (if r or p is revealed: $R_c \leftarrow R_c \setminus r$ or $P_c \leftarrow P_c \setminus p$).

We formalize PACCART's concealment component by providing it the ability to adopt a **privacy behavior**, consisting of three concealing aspects: **Scope**, **Division** and **Dedication**.

Scope: At each point in the dispute, if possible, an agent extends the dispute by adding one or more arguments (Algorithm 1, Step 4). The amount of useful arguments (as defined in Section 2.1) that an agent considers to add at any point of time to the dispute, is called its scope. An agent without any focused scope would add all available useful arguments at once. An agent with a focused scope is able to carefully select a smaller set of arguments, and locally gains control over the amount of the added (and therefore revealed) content. The larger the scope of an agent, the more content is added at each step in the dispute.

Division: Not all information is equally important. To be able to denote this, we split the sets of contents into **set-families** [4] of content. These subgroups can then be ordered to the likings of the agent. This entails splitting the knowledge base into ordered subgroups of different groups of conceal-worthy content. Therefore, based on the original knowledge base $KB = \{P, R, C\}$, we propose an **ordered subdivided knowledge base (OSKB)**, which includes the following ordered tuples of set-families:

- An ordered tuple of premises $O_P = \langle P_1, \dots, P_n \rangle$
- An ordered tuple of rules $O_R = \langle R_1, \dots, R_n \rangle$

The relation between these ordered set-families F_X and the sets X (with $X = P, R$) all follow the same properties:

- $\bigcup F_X = X$
- $\bigcap F_X = \emptyset$
- $\forall y \in Y, \forall z \in Z ((Y \subseteq F_X \wedge Z \subseteq F_X) \rightarrow (y = z \leftrightarrow Y = Z))$

With the introduced *OSKB*, an agent can order their content based on its concealment preferences. We can therefore treat these two ordered tuples together as one totally ordered knowledge base, subdivided in what we call **dedication levels**, as follows: $L = \langle \{O_{P_1}, O_{R_1}\}, \dots, \{O_{P_n}, O_{R_n}\} \rangle$. Each level contains one or more premises and rules. The first level L_1 contains content at the top of the ordering of each of the *OSKB* tuples, which is the content that the agent is the least concerned about revealing. The last level L_n contains content at the bottom of the ordering, indicating the content that the agent considers most important to conceal and therefore has to fully commit to winning the dispute in order to be willing to reveal these pieces of information. The **exhaustion** of

an agent’s division aspect indicates the amount ordered subdivisions an agent makes. The more exhaustive content subdivision, the higher amount of levels an agent splits its *OSKB* up into.

An example of four different *OSKB* divisions is shown in Figure 1, where an agent makes no subdivision of its *OSKB* (1a), by adding all its arguments to Level 1. Another possibility is it divides its *OSKB* in half, with two levels (1b). Furthermore, an agent can choose to divide its *OSKB* in all separate arguments, which yields four levels in this case (1c). Note that with this approach, Figure 1 shows an example of a level with just one premise (‘j’), as only one premise can suffice to form an argument. A final approach consists of an agent dividing its *OSKB* by subdividing all of its content (all rules and premises) over different levels, yielding ten levels in this case (1d).

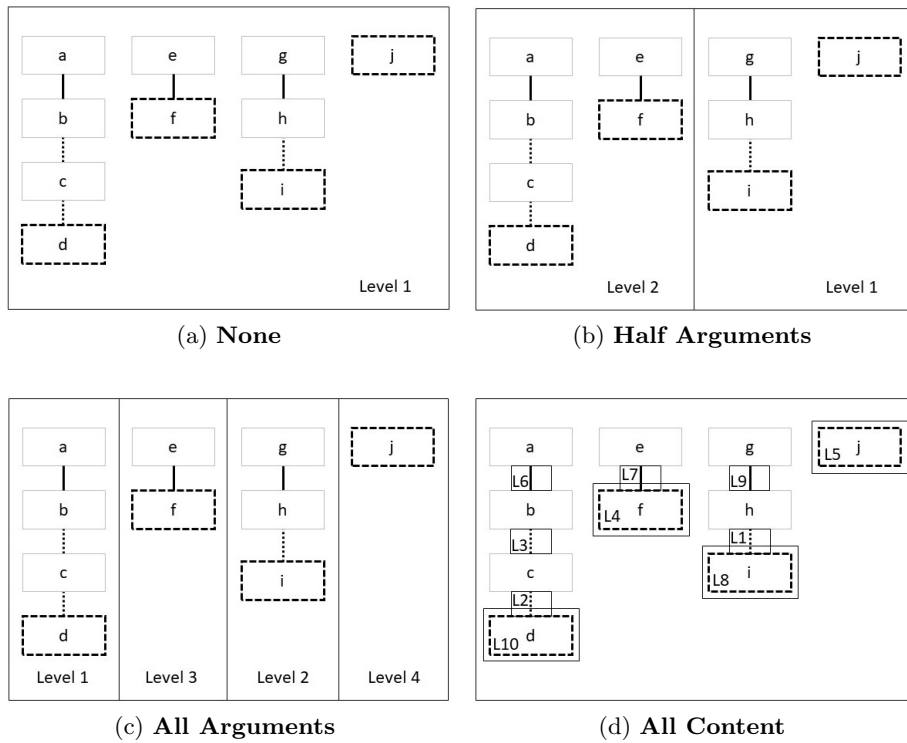


Fig. 1: Examples of different approaches of the PACCART agent’s division aspect. Four arguments consisting of ten pieces of content are divided up into different levels. Solid and dashed lines are for strict rules R_s and defeasible rules R_d , respectively.

Dedication: Agents that are able to divide their content into levels, can use this to their advantage. Such agents will initially only provide arguments if they can do so from their first level in their knowledge base. When all arguments in a first level have been depleted, the agents receive the option to either drop to a new level, therefore making a further argument privacy concession, or to forfeit the dispute. This gives agents the ability to weigh their decision to further dedicate to the argumentation. The amount of **willingness** to drop determines the agent’s dedication to continue the dispute. The more willing an agent is to drop dedication levels, the more it will use and therefore reveal the contents of its *OSKB*. This is calculated by whether a certain willingness Threshold θX with $X \in [0, 100]$ is met at the time of decision whether to commit further to the dispute. This means that an agent with $\theta 75$ has a 75% chance of dropping each level. This entails that the agent example of Figure 1c has a $0.75^3 = 42.2\%$ chance to use the content of its final level (as it could drop three times until it reaches its fourth and final level of content), whereas agent example of Figure 1d has a $0.75^9 = 7.5\%$ chance to fully commit its *OSKB*.

Any combination of all three concealing aspects maps to an agent’s **privacy type**. These privacy behaviors are in place for agents to further gain control over their content concealment during disputes, as well as influence their win rate.

3.2 Equity Component

Recall that we want our PACCART agent to be able to help different types of users to deliver on the equity aspect. On user’s privacy stances, we follow Dupree et al. [8], who determine a categorization based on stances regarding privacy along two dimensions. We define a user u with **knowledge** $k \in \{low, medium, high\}$ and **motivation** $m \in \{low, medium, high\}$. The degree of knowledge indicates the amount of awareness a user has about their privacy and the degree of general knowledge on privacy matters. The degree of motivation indicates the effort a user expends to protect their privacy and the degree of willingness to act on privacy matters. Each system user falls in one of five categories, also known as **privacy types**:

- **Fundamentalists**: high knowledge, high motivation
- **Lazy Experts**: high knowledge, low motivation
- **Technicians**: medium knowledge, high motivation
- **Amateurs**: medium knowledge, medium motivation
- **Marginally Concerned**: low knowledge, low motivation

Dupree et al. determine the rate at which users fall into these categories: 3% of users are Fundamentalists, Lazy Experts 22%, Technicians 18%, Amateurs 34% and Marginally Concerned 23%. This is comparable to the categorical distributions of privacy types of earlier conducted researches [1,6,8,18,25,29].

We define PACCART agents that adapt to the knowledge and motivation of the users’ privacy type as **Personalized** agents, whereas we consider **indifferent** agents to be not personalized and therefore have an unfocused scope and

make no distinction between the importance of content in their *KB*. In order for personalized agents to be considered equitable, they should adhere to the following equity properties, based on earlier research on equity [31,33]:

- EP1:** The knowledge and motivation of a user is considered and utilized to the fullest extent by their personalized agent.
- EP2:** A personalized agent outperforms an indifferent agent.
- EP3:** There are no performance outliers between personalized agents; no personalized agent heavily over- or underperforms compared to others.

EP1 is important because the strengths of the user should be taken into account by their agent. The privacy stance of a user should not be ignored, as this would be unfair towards users that are heavily engaged in protecting their privacy. In the same line, EP2 is important because the agents that are tailored towards a user should not perform worse than an agnostic, basic agent. Providing personalization should be beneficial for users, not disadvantageous. EP3 is important because in order to reach fair outcomes, it should not be the case that the privacy stance of a user exorbitantly influences the performance of their agent. It would e.g. be unfair towards unknowledgeable users if their agents would underperform by design.

In order to meet these properties, we introduce a mapping between users and agents. This way, both knowledge and motivation are used to determine the personalized agent’s privacy type. We determine a fitting mapping between users u to their agents a such that all users get mapped to the shortest scope, user knowledge is mapped to agent division and user motivation is mapped inversely to agent dedication. We will substantiate each mapping.

First, we assign all personalized agents to have a small scope, since a small scope is beneficial for all users, independent of privacy stance. When a user has a high privacy stance, they can let their agent subdivide its content in such a way that each piece of content is thoroughly protected. This would mean that the agent already has a small amount of content to choose from, so for a high privacy user the scope has only a little positive impact. However, for users who do not have a lot of knowledge or motivation to bring to the dispute, a small scope is also the best fit as it protects as much content as possible.

Secondly, we map a user’s knowledge to their agent’s division, because of the degree of user knowledge should correspond with the amount of useful subdivisions of their agent’s *OSKB* levels. This means that the higher the user’s knowledge, the higher the agent’s content dividing. Someone with a high knowledge could benefit from an agent with a high capability of dividing its knowledge base content. This would allow users to provide their agent with their preferences in detail. This is in line with EP1. Similarly, mapping a low knowledge to a low *OSKB* division would also be useful. This is because users with low knowledge have little relevant preference divisions to make in their agent’s knowledge base.

Thirdly, we map a user’s motivation inversely to their agent’s dedication, because the amount of motivation of a user should correspond to the dedication of its agent to conceal content (in favor of winning disputes). This means that the higher the user’s motivation, the lower the agent’s dedication. Users that are

Table 1: All three concealing aspects of indifferent PACCART agent and personalized PACCART agents that are matched with representative agents for different user privacy types.

Privacy Type	Scope	Division	Dedication
<i>Indifferent</i>	All	None	$\theta 100$
Fundamentalist	Shortest	AllContent	$\theta 25$
Technician	Shortest	AllArgs	$\theta 25$
Amateur	Shortest	AllArgs	$\theta 50$
Lazy Expert	Shortest	AllContent	$\theta 75$
M.Concerned	Shortest	HalfArgs	$\theta 75$

highly motivated to protect their data would rather have their agent drop as little levels as possible, even if it would require taking (social) losses. Similarly, users that prefer not to act on privacy matters would want their agents to perform well when it comes to winning disputes, but would not mind agents revealing information to do so. This is also in line with EP1. This mapping results in five personalized agents, one representative for each user type, as noted in Table 1. This table also includes an indifferent agent.

3.3 Additional Usability Components

In addition to the Concealment and Equity components two usability measures are taken. A Collaboration component is introduced to support both sides of the dispute to be represented by multiple agents. This is achieved by introducing the notion of teams such that the set of agents A in the protocol now consists of $A = \{T_p, T_o\}$ to support both a proponent team $T_p = \{a_{p1}, \dots, a_{pn}\}$ and opponent team $T_o = \{a_{o1}, \dots, a_{on}\}$. In order to extend a dispute each team of PACCART agents continuously selects one of its agents to extend. A team forfeits when none of its agents can extend the dispute any further. This component allows for multiple PACCART agents to cooperate on a common goal of defending/attacking a privacy related subject. This means that agents can add content from their own *OSKB* to the dispute when other agents in their team fail to do so.

Furthermore, an Explainability component is introduced to give users insights to the working of their agent. The semantic nature of PACCART allows us to produce both textual and visual output. PACCART can provide textual output by considering outcomes and providing feedback to the user. Based on this, it is able to give different kinds of feedback, with a range of detail. It can notify users on a summary (e.g., *"I have won 56% of today's disputes and managed to conceal 73% of your content"*) or it can give detailed advice on possible actions to be taken to improve its performance (such as listing possible weak points in its arguments for the user to improve upon). Furthermore, PACCART can provide visual output by showing its user images of the Structured Argumentation Framework [22] of final disputes. This gives users a visual overview

of (counter)arguments and possible weak points in their content. This component allows users of PACCART to better understand its inner workings and performance.

4 Experimental Results

The PACCART agent and the experimental setup are implemented as a C# program. For the sake of reproducibility, we make this program and experiments open source, along with examples and schematic overviews of the PACCART agent workings.³

4.1 Dataset Generation

We implement a system that generates datasets of disputes according to four parameters. The **disputeAmount** parameter indicates the amount of generated unique disputes. A higher input value indicates a larger set of disputes, therefore less prone to outliers. The **disputeSize** parameter controls the amount of arguments that the dispute can contain. A higher input value indicates larger disputes with more content. The **maxArgumentSize** parameter dictates the maximum amount of subarguments that each argument can consist of. A higher input value indicates larger arguments with more content and therefore more attackable weak points. Finally, **maxBranches** is used to control the maximum amount of attacks that each weak point can have, indicating a branching choice in the dispute. A higher input value indicates more options for both agents.

By tuning these parameters, we are able to generate dispute datasets of various shapes and sizes, which makes for exhaustive possibilities for testing functionalities of PACCART. After preliminary analysis of variables, we generate a dispute dataset based on the default parameter settings ($\text{disputeAmount} = 200$, $\text{disputeSize} = 20$, $\text{maxArgumentSize} = 10$, $\text{maxBranches} = 2$).

4.2 Experiment 1: Effect of Privacy Behaviors

Setting The goal of the first experiment is to test the performance of PACCART agents. Agent performance is evaluated on two metrics, average concealment C_{avg} and average win rate W_{avg} . We hypothesize the following:

- H1:** A smaller scope leads to both increased concealment and increased win rate.
- H2:** More exhaustive division leads to increased concealment and decreased win rate.
- H3:** A higher dedication leads to decreased concealment and increased win rate.

We determine four or five conditions for each of the three privacy behavior aspects, to test the range of PACCART’s concealing behaviors. For the scope, we include selecting the **Shortest** or **Longest** arguments, as well as a **Random** argument or **All** possible arguments. For the division, we follow the examples of Figure 1 and include conditions where **None** of the content is split, where the

OSKB is split into two groups of arguments (**HalfArgs**), split into all separate levels of arguments (**AllArgs**) or a subdivision where each level contains a single piece of content (**AllContent**). The dedication conditions consist of an increasing threshold θ , with $\theta \in \{0, 25, 50, 75, 100\}$ that should be met in order to drop to a new level. These conditions yield 80 possible privacy types. Each of these 80 predetermined agents are set up against all other agents, and simulations are run on the 200 disputes of our dataset. This means that the experiment is run on 16,000 disputes for 80 agent set-ups, totaling in 1,280,000 simulated disputes. For each of the disputes, both agents are evaluated as a proponent, as well as opponent of the dispute, to ensure equal chances of winning.

Results Figures 2 and 3 depict the performance of the 80 different agent privacy behavior types, across all three concealing aspects.




Fig. 2: Average Concealment C_{avg} results for all PACCART privacy types.

Scope: We observe from Figures 2 and 3 that the scope of an agent has a significant effect on its performance. Both the average win rate W_{avg} and average concealment C_{avg} increase with a smaller scope. We conclude that a smaller scope has a strictly positive impact. This confirms hypothesis H1.

Dividing: All of the **None** dividing aspect results are equal, independent of dropping willingness. This means that not dividing the *OSKB* negates the effect of the agent’s dedication. This is an expected outcome, which happens because there is no division made of the knowledge base so there are no levels for the agent to drop between, even if it would be willing. Outside of this behavior, an upward trend is noticeable in all cases for average concealment, as well as a downward trend in all cases for the win rate, with more exhaustive dividing. This confirms hypothesis H2.

Dedication: When looking at the dedication aspect, we observe an upward trend in all cases for average concealment C_{avg} , as well as a downward trend in

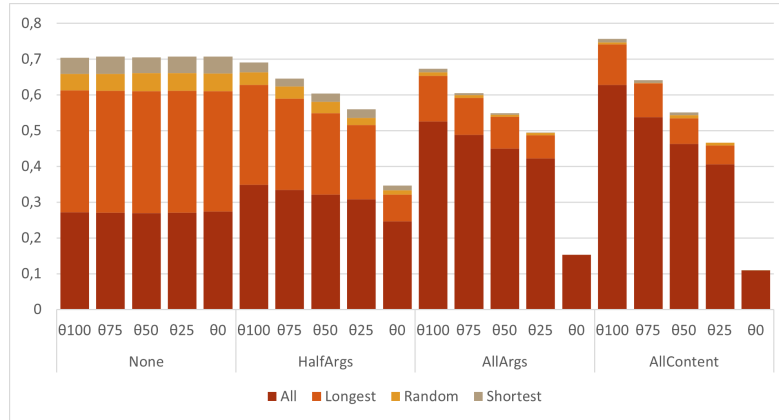


Fig. 3: Average Win Rate W_{avg} results for all PACCART privacy types.

all cases for win rate W_{avg} , with less willing dedication. This is a similar trend as with the dividing aspect of the privacy behavior. This confirms hypothesis H3. Furthermore, Figure 3 shows a significant drop in win rate from θ_{25} to θ_0 , while the improvement in concealment is disproportional. This shows that it is beneficial for an agent to be at least somewhat willing to commit to the dispute. Based on these results, we conclude the following observation:

Observation 1. *PACCART's concealment component allows users to keep information private, while also giving them the choice of a trade-off between winning disputes and further protection of information.*

4.3 Experiment 2: Effect of User-Agent Mapping in Realistic Setting

Setting The goal of the second experiment is to evaluate the mappings between agents and users by simulating disputes for each personalized agent in a realistic setting. The results of this mapping will determine whether EP2 and EP3 are met, which means that PACCART is an equitable agent. Therefore, based on this mapping, we further hypothesize:

- H4:** Equity property EP2 is met under a mapping where personalized agents are assigned the smallest possible scope.
- H5:** Equity property EP3 is met under a mapping where personalized agents are assigned a fitting trade-off between division and dedication.

We create a set of opponents according to data of distribution of real life user population as given by Dupree et al. This opponent set therefore contains three Fundamentalist agents, 22 Lazy Expert agents, 18 Technician agents, 34 Amateur agents and 23 Marginally Concerned agents. We call this set of 100 agents the **Model Population Set MPS**. The MPS is in place because in a

practical scenario it is less likely that an MPC occurs between Fundamentalists’ agents, as between Marginally Concerned users’ agents.⁴

This means that six agents (one indifferent agent and all five personalized agents) compete 100 times against each of the personalized agents, and simulations are run on 200 disputes on the dispute dataset. Overall, the experiment is run on 20.000 disputes for six agent set-ups. Furthermore, agents are again tested twice for all disputes, both as proponent and opponent of the subject, to ensure equal chances of winning.

Results The results of the second experiment can be seen in Figure 4. Again, performance is measured by concealment C_{avg} and win rate W_{avg} . As shown in Figure 4, the indifferent agent performs much worse than the personalized agents on both metrics (only 0.185 for win rate and 0.660 for concealment). This confirms hypothesis H4.

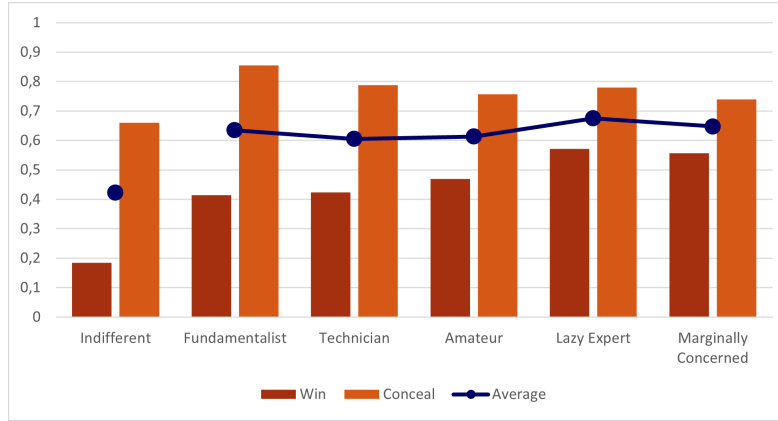


Fig. 4: Average win rate W_{avg} and Average Concealment C_{avg} for indifferent agent and personalized agents in MPS. Averages between W_{avg} and C_{avg} are indicated with a line.

Furthermore, the averages of all personalized agents range between 0.6 and 0.7. This means that although some personalized agents are better at winning or concealing, the overall performance leads to an equitable situation where no users are victimized by the agent’s workings. This confirms hypothesis H5. It is worth noting that an interesting trend occurs between personalized agents, where the Fundamentalist representative’s agent (with the highest privacy stance) wins the least and conceals the most, while the Marginally Concerned representative’s agent wins the most and conceals the least. This trade-off shows how the different

⁴ An additional experiment is performed to evaluate the MPS, placing all agents in a non-distributed setting, which yielded similar results.

privacy stances influence the results. Based on these results, we conclude the following observation:

Observation 2. *PACCART’s equity component allows for a well-matched personalization for users of various privacy stances. While personalized PACCART agents overall perform relatively well, a consistent trade-off between win rate and concealment shows that no user is disadvantaged.*

5 User Study

We further conduct a user study to understand what components of PACCART lead to user trust.

5.1 Setting

We design a survey in two parts. The first part of the survey has questions on the privacy stance of participants, in order to assess their privacy type. We deliberately use existing questions from the literature to ensure compatibility: three questions used by Westin et al. (e.g., *“How much do you agree with the statement ‘Most businesses handle the personal information they collect about consumers in a proper and confidential way.’?”*) [15] to determine the knowledge of participants on privacy and 10 questions on statements about privacy from the study of Dupree et al., to determine the motivation of participants on privacy (e.g., *“How strongly do you identify yourself with the statement ‘I would rather choose being social over privacy.’?”*). As validation and to mitigate response bias, we also ask participants directly to self-assess their own knowledge and motivation (e.g., *“How much do you know about digital privacy issues?”*). These questions are all answered on a Likert scale. The full questionnaire is also made openly available.³

The second part of the survey has questions on the various components of PACCART as a personal assistant. This part starts with an example scenario. Then a set of questions follows in which participants are asked to rate their perceived trust of such personal assistants on a Likert scale (1 = Strongly Distrust, 5 = Strongly Trust). The first question is on the participants’ initial thoughts of trust on the PACCART base component (an explanation followed by *“How much would you trust to use such a privacy assistant?”*). Then, each separate PACCART component is explained separately and addressed as a question. Afterwards, the participants are asked to rate the agent with all components combined (the base component with all four additional components). Finally, the participants are asked to reconsider their thoughts on the base component. This gives the participants a chance to reflect on their initial thoughts.

The survey is distributed through Qualtrics, an online, secure cloud-based, survey tool. Data is automatically and anonymously recorded through Qualtrics, in accordance with GDPR requirements. The survey is preceded by filling out a consent form. To ensure correctness and clarity, we first perform a small pilot study. Afterwards, the survey is distributed online for the user study.

5.2 Results

Data was collected from 117 voluntary participants in the user study. Based on validation questions and completion requirements, 12 survey responses are filtered out. Out of the remaining 105 participants, eight participants self-assessed as Fundamentalists, 20 participants as Lazy Experts, 22 as Technicians, 31 as Amateurs and 24 as Marginally Concerned users. This is in line with the distributions by Dupree et al. [8].

We report the mean (M) and standard deviation (SD) of the results, as well as significance through t-tests (P). The results indicate that the initial consideration of the PACCART base component is fairly neutral (M = 2.857, SD = 1.023), slightly leaning towards distrust. The trust ratings given by participants are higher than the initial consideration for both Concealment (M = 2.943, SD = 0.979) as well as Equity (M = 3.171, SD = 1.069). There is a significantly ($P < .001$) positive increase of trust of the combined agent (M = 3.467, SD = 0.974) compared to the initial consideration of the base component. Even more so, when asked to reevaluate the trustworthiness of the agent, the average trust rating significantly ($P < .001$) drops (M = 2.362, SD = 0.982) compared to the combined agent.




Fig. 5: Average survey study ratings of participant trust on a Likert scale (1 = Strongly Distrust, 2 = Distrust, 3 = Neutral, 4 = Trust, 5 = Strongly Trust), divided by privacy type.

The results are split on each of the privacy type categories, as shown in Figure 5. The graph shows the average trust ratings by users with different privacy types. From this graph, we can observe the effect of different stances on privacy on agent trust scores. These results show that the lower the privacy stance, the higher the overall trust in the agent. Fundamentalist participants' highest mean trust rating is a 2.130, whereas Marginally Concerned participants'

lowest mean trust rating is a 2.727. This is in line with our expectations about the privacy types and therefore an indication that the privacy stance assessment part of the survey works as intended.

A further noteworthy observation is that for all of the individual privacy types the reconsideration is rated lower than the initial thoughts on the base component. This indicates that after having read an explanation on what possible components could improve upon the base, participants independently of their privacy type assess the base component to be less trustworthy. When comparing the base component with the total combined agent, trust significantly increases for all user types ($P < .001$) except for Fundamentalists. While the results do indicate an increase of trust for Fundamentalists, the results are not significant ($P = .18$), which is expected because of the naturally low occurrence of users with this high privacy stance. These results strongly indicate that overall, the principles of PACCART and its components increase the indicated trust of users of all privacy stances.

6 Conclusion

We introduced PACCART, which helps users preserve privacy by enabling automated privacy argumentation. PACCART aims to induce trust by increasing content concealment, providing equitable personalizations, enabling multiagent team-based collaboration and explaining its actions through feedback. The agent is designed to be general and is made publicly available as an open-source program together with the dispute dataset generation system, so that they can be used for research as well as in practical applications, such as team collaboration tools (e.g., MS Teams) where co-owned data is shared abundantly and privacy disputes need to be resolved.

Future research could further investigate what improvements the system needs for its proper use in open systems, for example setting up standardization of the use of *OSKB*'s and determining what information is shared beforehand to increase privacy for both parties involved. Another research avenue would be to close the feedback loop between users and the agent to further increase trust. When users get prompted that their agent lost a dispute because of the lack of arguments, the user could respond by taking action to help and improve the agent fit to its user. Furthermore, introducing mutual feedback opens new possibilities for machine learning approaches. Now, there exists a mapping between users and their personalized agents, which could be changed into the agent learning the preferences of the user instead. Weights could be given to the importance of dedication to win certain disputes, or concealing specific levels of content. The inclusion of reinforcement learning could be an important additional step towards robust and well-adjusted argumentation based privacy assistants.

References

1. Ackerman, M.S., Cranor, L.F., Reagle, J.: Privacy in e-commerce: examining user scenarios and privacy preferences. In: Proceedings of the 1st ACM Conference on Electronic Commerce. pp. 1–8 (1999)
2. Addo, I.D., Ahamed, S.I., Yau, S.S., Buduru, A.: A reference architecture for improving security and privacy in internet of things applications. In: IEEE International Conference on Mobile Services. pp. 108–115. IEEE (2014)
3. Baroni, P., Caminada, M., Giacomin, M.: An introduction to argumentation semantics. *The knowledge engineering review* **26**(4), 365–410 (2011)
4. Brualdi, R.A.: *Introductory combinatorics*. Pearson Education India (1977)
5. Colnago, J., Feng, Y., Palanivel, T., Pearman, S., Ung, M., Acquisti, A., Cranor, L.F., Sadeh, N.: 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. pp. 1–13 (2020)
6. Consolvo, S., Smith, I.E., Matthews, T., LaMarca, A., Tabert, J., Powledge, P.: Location disclosure to social relations: why, when, & what people want to share. In: Proceedings of the SIGCHI conference on Human factors in computing systems. pp. 81–90 (2005)
7. Cyras, K., Toni, F.: ABA+: assumption-based argumentation with preferences. In: Fifteenth International Conference on the Principles of Knowledge Representation and Reasoning. pp. 553–556 (2016)
8. Dupree, J.L., Devries, R., Berry, D.M., Lank, E.: Privacy personas: Clustering users via attitudes and behaviors toward security practices. In: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems. pp. 5228–5239 (2016)
9. Fazzinga, B., Galassi, A., Torroni, P.: An argumentative dialogue system for covid-19 vaccine information. In: International Conference on Logic and Argumentation. pp. 477–485. Springer (2021)
10. Fogues, R.L., Murukannaiah, P.K., Such, J.M., Singh, M.P.: 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)
11. Fogues, R.L., Murukannaiah, P.K., Such, J.M., Singh, M.P.: Sosharp: Recommending sharing policies in multiuser privacy scenarios. *IEEE Internet Computing* **21**(6), 28–36 (2017)
12. Jin, H., Guo, B., Roychoudhury, R., Yao, Y., Kumar, S., Agarwal, Y., Hong, J.I.: Exploring the needs of users for supporting privacy-protective behaviors in smart homes. In: CHI Conference on Human Factors in Computing Systems. pp. 1–19 (2022)
13. Kökciyan, N., Yaglikci, N., Yolum, P.: An argumentation approach for resolving privacy disputes in online social networks. *ACM Transactions on Internet Technology* **17**(3), 1–22 (2017)
14. Kökciyan, N., Yolum, P.: Priguard: A semantic approach to detect privacy violations in online social networks. *IEEE Transactions on Knowledge and Data Engineering* **28**(10), 2724–2737 (2016)
15. Kumaraguru, P., Cranor, L.F.: Privacy indexes: a survey of Westin’s studies. Carnegie Mellon University, School of Computer Science, Institute for Software Research International (2005)

16. Kurtan, A.C., Yolum, P.: Assisting humans in privacy management: an agent-based approach. *Autonomous Agents and Multi-Agent Systems* **35**(1), 1–33 (2021)
17. Maple, C.: Security and privacy in the internet of things. *Journal of Cyber Policy* **2**(2), 155–184 (2017)
18. Maus, B., Olsson, C.M., Salvi, D.: Privacy personas for iot-based health research: A privacy calculus approach. *Frontiers in digital health* p. 187 (2021)
19. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., Galstyan, A.: A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* **54**(6), 1–35 (2021)
20. Mester, Y., Kökciyan, N., Yolum, P.: Negotiating privacy constraints in online social networks. In: *International Workshop on Multiagent Foundations of Social Computing*. pp. 112–129. Springer (2015)
21. Miller, T.: Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence* **267**, 1–38 (2019)
22. Modgil, S., Prakken, H.: The aspic+ framework for structured argumentation: a tutorial. *Argument & Computation* **5**(1), 31–62 (2014)
23. Mosca, F., Such, J.M., McBurney, P.: Towards a value-driven explainable agent for collective privacy. In: *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems*. pp. 1937–1939 (2020)
24. O’neil, C.: *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown (2017)
25. Sheehan, K.B.: Toward a typology of internet users and online privacy concerns. *The information society* **18**(1), 21–32 (2002)
26. Squicciarini, A.C., Novelli, A., Lin, D., Caragea, C., Zhong, H.: From tag to protect: A tag-driven policy recommender system for image sharing. In: *2017 15th Annual Conference on Privacy, Security and Trust (PST)*. pp. 337–33709. IEEE (2017)
27. Story, P., Smullen, D., Yao, Y., Acquisti, A., Cranor, L.F., Sadeh, N., Schaub, F.: Awareness, adoption, and misconceptions of web privacy tools. *Proceedings on Privacy Enhancing Technologies* **2021**(3), 308–333 (2021)
28. Such, J.M., Porter, J., Preibusch, S., Joinson, A.: Photo privacy conflicts in social media: A large-scale empirical study. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. pp. 3821–3832 (2017)
29. Taylor, H.: Most people are “privacy pragmatists” who, while concerned about privacy, will sometimes trade it off for other benefits. *The Harris Poll* **17**(19), 44 (2003)
30. Tonge, A., Caragea, C.: Image privacy prediction using deep neural networks. *ACM Transactions on the Web (TWEB)* **14**(2), 1–32 (2020)
31. Ulusoy, O., Yolum, P.: Panola: A personal assistant for supporting users in preserving privacy. *ACM Transactions on Internet Technology* **22**(1), 1–32 (2021)
32. Wishart, R., Corapi, D., Marinovic, S., Sloman, M.: Collaborative privacy policy authoring in a social networking context. In: *2010 IEEE International Symposium on Policies for Distributed Systems and Networks*. pp. 1–8. IEEE (2010)
33. Woodgate, J., Ajmeri, N.: Macro ethics for governing equitable sociotechnical systems. In: *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*. pp. 1824–1828 (2022)