

PACCART: Reinforcing Trust in Multiuser Privacy Agreement Systems

Extended Abstract

Daan Di Scala
Utrecht University
Utrecht, Netherlands
daandiscala@hotmail.com

Pinar Yolum
Utrecht University
Utrecht, Netherlands
p.yolum@uu.nl

ABSTRACT

Content in collaborative systems, such as Online Social Networks, is often co-owned by multiple users with different privacy expectations, leading to possible multiuser privacy conflicts. In order to address these conflicts, 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.

KEYWORDS

Multiuser Privacy; Trust; Equity

ACM Reference Format:

Daan Di Scala and Pinar Yolum. 2023. PACCART: Reinforcing Trust in Multiuser Privacy Agreement Systems: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 3 pages.

1 INTRODUCTION

Privacy is the right of individuals to keep personal information to themselves [22]. Recent research on managing privacy online shows promising results on how agents can help with privacy, such as on detecting privacy violations [10], recommending sharing behavior [8, 18], and learning privacy preferences [11, 21]. An interesting problem of privacy arises in collaborative systems, such as Online Social Networks, where the content being shared is generally **co-owned**, 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 [7]). 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)** [16, 20].

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 and acts 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 [9, 19]. Many existing studies of collaborative systems indicate the importance of *trust* in making systems usable by individuals [3, 12]. 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. Therefore, users would know that their privacy is safe with the agent [1, 12].

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 [13, 17], 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 [22, 24].

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 [23]. 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 [19]. 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 [6, 14].

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. PACCART is openly available at: <https://github.com/PACCART/PACCARTpaper>.

2 MODEL

The PACCART agent consists of a base component, which allows it to communicate through a dialogical argumentation framework. Four additional components are introduced on top of the workings of the base component.

We formalize PACCART’s concealment component by providing it the ability to adopt a **privacy behavior**, consisting of a combination of three concealing aspects: **scope**, **division** and **dedication**.

Scope: Agents can often choose between many useful arguments. The amount of useful arguments that an agent considers at any point is called its scope. Agents with an unlimited scope reveal all available useful arguments at once. With its adjustable scope, PACCART is able to carefully select a smaller set of arguments, and therefore locally gains control over the amount of revealed content.

Division: Not all information is equally important. Agents with unordered knowledge bases have no control over which arguments to prioritize over others. Therefore, we propose an **ordered subdivided knowledge base (OSKB)**, which entails splitting the knowledge base into ordered subgroups (or set-families [2]) of different groups of conceal-worthy content. With the introduced OSKB, PACCART can order its content based on its concealment preferences and prioritize which information to withhold.

Dedication: Agents can be conservative with their content by only providing arguments from their least important set of arguments in their OSKB. However, these initial arguments might not be enough to win a dispute. PACCART is able to adjust its level of dedication to reveal more ordered content in order to try to win more disputes. When all initial arguments have been used, the agent has the option to either withhold the rest of its content and forfeit the dispute or to start revealing more important arguments and therefore make further privacy concessions. This gives agents the ability to weigh their decision to further dedicate to the argumentation. The more dedicated the agent is, the more OSKB content it reveals.

To deliver on the Equity aspect, we want PACCART agent to be able to help different types of users. On user’s privacy stances, we follow Dupree et al. [5], 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 user falls in one of five categories, also known as **privacy types**.

In order for PACCART to be an equitable agent, it should assist users of all privacy types, so all users benefit equally from their agent. To achieve this, we introduce **personalized** agents for each of the five privacy types, by mapping the three concealment aspects accordingly to the degrees of k and m . **Indifferent** agents are agents that are not personalized and thus have an unfocused scope and make no distinction between the importance of content in their KB.

We introduce a Collaboration component to support both sides of the dispute to be represented by multiple agents. 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.

Finally, we introduce an Explainability component 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 with a summary or give detailed advice on possible actions to be taken to improve its performance. Furthermore, PACCART can provide visual output by showing its user images of the Structured Argumentation Framework [15] 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.

3 MAIN RESULTS

We have conducted experiments to evaluate the performance of PACCART on concealment and equity. The PACCART agent implementation and experimental setup are made open source. Due to the limited availability of datasets of argumentation, we developed a system that generates datasets of disputes according to four parameters. By tuning these parameters, we are able to generate dispute datasets of various shapes and sizes, which makes for exhaustive possibilities for testing the performance of PACCART.

Based on the simulations that we have ran we observe that 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. A smaller scope strictly increases both concealment and win rate of the agent. However, when the agent deploys a more exhaustive OSKB division and a lower dispute dedication, it positively impacts its concealment and negatively impacts its win rate.

Furthermore, we observe that PACCART’s equity component allows for a well-matched personalization for users of all privacy stances. Personalized PACCART agents overall perform well and a consistent trade-off between win rate and concealment shows that no user type is disadvantaged. Finally, personalized agents outperform indifferent (non-personalized) agents consistently on both concealment and win rate, affirming that personalizing the agent is beneficial to all users.

4 CONCLUSION

We introduced PACCART, which helps users preserve their 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. The complete details regarding PACCART’s implementation, evaluation through comparison with relevant related works and a user study showcasing the effect of perceived trust of PACCART are presented in [4].

5 ACKNOWLEDGMENTS

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

REFERENCES

- [1] Ivor D Addo, Sheikh I Ahamed, Stephen S Yau, and Arun Buduru. 2014. A reference architecture for improving security and privacy in internet of things applications. In *IEEE International Conference on Mobile Services*. IEEE, 108–115.
- [2] Richard A Brualdi. 1977. *Introductory combinatorics*. Pearson Education India.
- [3] Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2020. Informing the design of a personalized privacy assistant for the internet of things. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [4] Daan Di Scala and Pinar Yolum. 2023. PACCART: Reinforcing Trust in Multiuser Privacy Agreement Systems. <https://doi.org/10.48550/ARXIV.2302.13650>
- [5] Janna Lynn Dupree, Richard Devries, Daniel M Berry, and Edward Lank. 2016. Privacy personas: Clustering users via attitudes and behaviors toward security practices. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 5228–5239.
- [6] Bettina Fazzinga, Andrea Galassi, and Paolo Torroni. 2021. An argumentative dialogue system for COVID-19 vaccine information. In *International Conference on Logic and Argumentation*. Springer, 477–485.
- [7] Ricard L Fogues, Pradeep K Murukannaiah, Jose M Such, and Munindar P Singh. 2017. Sharing policies in multiuser privacy scenarios: Incorporating context, preferences, and arguments in decision making. *ACM Transactions on Computer-Human Interaction (TOCHI)* 24, 1 (2017), 1–29.
- [8] Ricard L Fogues, Pradeep K Murukannaiah, Jose M Such, and Munindar P Singh. 2017. Sosharp: Recommending sharing policies in multiuser privacy scenarios. *IEEE Internet Computing* 21, 6 (2017), 28–36.
- [9] Haojian Jin, Boyuan Guo, Rituparna Roychoudhury, Yaxing Yao, Swarun Kumar, Yuvraj Agarwal, and Jason I Hong. 2022. Exploring the Needs of Users for Supporting Privacy-Protective Behaviors in Smart Homes. In *CHI Conference on Human Factors in Computing Systems*. 1–19.
- [10] Nadin Kökciyan and Pinar Yolum. 2016. Priguard: A semantic approach to detect privacy violations in online social networks. *IEEE Transactions on Knowledge and Data Engineering* 28, 10 (2016), 2724–2737.
- [11] Abdurrahman Can Kurtan and Pinar Yolum. 2021. Assisting humans in privacy management: an agent-based approach. *Autonomous Agents and Multi-Agent Systems* 35, 1 (2021), 1–33.
- [12] Carsten Maple. 2017. Security and privacy in the internet of things. *Journal of Cyber Policy* 2, 2 (2017), 155–184.
- [13] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* 54, 6 (2021), 1–35.
- [14] Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence* 267 (2019), 1–38.
- [15] Sanjay Modgil and Henry Prakken. 2014. The ASPIC+ framework for structured argumentation: a tutorial. *Argument & Computation* 5, 1 (2014), 31–62.
- [16] Francesca Mosca, Jose M Such, and Peter McBurney. 2020. Towards a value-driven explainable agent for collective privacy. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems*. 1937–1939.
- [17] Cathy O’neil. 2017. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.
- [18] Anna Cinzia Squicciarini, Andrea Novelli, Dan Lin, Cornelia Caragea, and Haoti Zhong. 2017. From tag to protect: A tag-driven policy recommender system for image sharing. In *2017 15th Annual Conference on Privacy, Security and Trust (PST)*. IEEE, 337–33709.
- [19] Peter Story, Daniel Smullen, Yaxing Yao, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2021. Awareness, adoption, and misconceptions of web privacy tools. *Proceedings on Privacy Enhancing Technologies* 2021, 3 (2021), 308–333.
- [20] Jose M Such, Joel Porter, Sören Preibusch, and Adam Joinson. 2017. Photo privacy conflicts in social media: A large-scale empirical study. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 3821–3832.
- [21] Ashwini Tonge and Cornelia Caragea. 2020. Image privacy prediction using deep neural networks. *ACM Transactions on the Web (TWEB)* 14, 2 (2020), 1–32.
- [22] Onuralp Ulusoy and Pinar Yolum. 2021. PANOLA: A Personal Assistant for Supporting Users in Preserving Privacy. *ACM Transactions on Internet Technology* 22, 1 (2021), 1–32.
- [23] Ryan Wishart, Domenico Corapi, Srdjan Marinovic, and Morris Sloman. 2010. Collaborative privacy policy authoring in a social networking context. In *2010 IEEE International Symposium on Policies for Distributed Systems and Networks*. IEEE, 1–8.
- [24] Jessica Woodgate and Nirav Ajmeri. 2022. Macro Ethics for Governing Equitable Sociotechnical Systems. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*. 1824–1828.