

Privacy-Aware Explanations for Team Formation^{*}

Athina Georgara^{1,2}, Juan A. Rodríguez-Aguilar¹, and Carles Sierra¹

¹ Artificial Intelligence Research Institute, CSIC, Barcelona, Spain

² Enzyme Advising Group, Barcelona, Spain

{ ageorg, jar, sierra } @ iia.csic.es

Abstract. Over the recent years there is a growing move towards explainable AI (XAI). The widespread use of AI systems in a large variety of applications that support humans decisions leads to the imperative need for providing explanations regarding the AI systems functionality. That is, explanations are necessary for earning the users trust regarding the AI systems. At the same time, recent legislation such as GDPR regarding data privacy require that any attempt towards explainability shall not disclose private data and information to third-parties. In this work we focus on providing privacy-aware explanations in the realm of team formation scenarios. We propose the means to analyse whether an explanation leads an explainability algorithm to incur in privacy breaches when computing explanation for a user.

Keywords: Privacy Awareness · Explainable AI (XAI) · Team Formation · Explainable Multi-Agents Environments (xMASE)

1 Introduction

Over the past decades there is wide interest in using artificial intelligence (AI) to aid humans to carry out complex, hard, and time-consuming tasks. As AI systems pervade our lives, people are becoming curious regarding the rationale and the methodology of these systems; thus we observe a new surge of interest towards *explainable AI (XAI)* [10, 22]. XAI provides “inside information” regarding the inner functionality of an AI system in an attempt to be transparent, and earn in this way the users’ trust. More and more applications turn to AI in order to ease and automate complex procedures, and demand understanding the solutions recommended by such systems. Besides the growing need for explanations, Goodman and Flaxman [15] point out that legislation such as the GDPR recently put forward by the EU leads to the *right to explanation*. That is, a user providing personal information as input data to some AI algorithm, has the right to know why the algorithm makes a decision with their input data instead of another one.

^{*} Research supported by projects AI4EU (H2020-825619), TAILOR (H2020-952215), 2019DI17, Humane-AI-Net (H2020-952026), Crowd4SDG (H2020-872944), and grant PID2019-104156GB-I00 funded by MCIN/AEI/10.13039/501100011033.

Doilovi et al. in [10] thoroughly discuss the intererpretability and the explainability of supervised machine learning models. The main focus of recent literature in explainability lies on machine learning (ML) models [1, 8, 10, 16, 23, 28, 29] and recommender systems [4, 17, 20, 21, 24, 33], which are usually considered as ‘black-boxes’, and transparency is a necessity. Beyond ML, Borg and Bex [6] recently introduced a general framework to provide contrastive explanations for (abstract) argumentation-based conclusions; Nardi et al. in [25] and Boixel and Endriss in [5] developed algorithms to justify outcomes (i.e., winners) in voting settings; and Georgara et al. in [13] propose a general algorithm to deliver explanations in team formation scenarios.

Recently, Kraus et al. [19] have raised awareness on the need for explanations in multiagent environments (xMASE), and they have identified the key challenges towards xMASE. Among other challenges—such as the development of appropriate algorithms for generating explanations and the user modelling to appropriately tailor explanations and increase user satisfaction—Kraus et al. refer to the issue of non-disclosing private data and information. Note that in any AI system that assists people in making a decision or solve a problem, individuals need to feed the system with information (possibly private), which is therefore utilised by the system to reach a solution. As such, within environments where explanations need to justify solutions involving many individuals, it is of utmost importance to ensure that private information remain private.

Need for privacy-awareness has risen as more and more data become available to AI systems. Considering the online social networks, we find privacy issues as people may expose data not only about themselves but also about others (e.g., via pictures, check-ins, etc.). Such and Criado in [32] discuss the multi-party privacy problem on online social networks, and highlight the need for mechanism that preven privacy violations in such environments. [18] works towards privacy in social networks, and develops a tool for detecting privacy violations in such settings using an agent-based representation for social networks. In a different to social media domain, Sörries et al. in [30] study privacy preserving technologies by design within the domain of healthcare. Now in XAI, Puiu et al. in [26] present recent developments on explainability and interpretability along with the limitation of data accessibility due to ethical constraints in cardiovascular diagnosis; while Sorvano et al. in [31] make a separation between *explainable* (X-) and *explanatory* (Y-) AI, and propose a model for YAI under GDPR guidelines.

In this paper we address the challenge of preserving privacy upon providing explanation within multi-agent environments, and specifically in team formation scenarios. Specifically, we argue that an AI system should only offer explanations that are guaranteed not to breach privacy. To the best of our knowledge this is the first work tackling this challenge in team formation. As such here we propose a *privacy breach detector* capable of finding whether a given explanation is bound to lead to privacy breaches. That is, we describe how our privacy breach detector interacts with a team formation algorithm (AI system) and an explanatory algorithm (XAI system) to approve or disapprove explanations within a general framework for privacy-aware explanations in team formation.

2 Background: Team Formation and Explanations

A team formation problem [2, 9, 3, 7, 12, 14] deals with situations where individuals must be grouped in teams to work on some task(s). In general, in such a problem there is a set of tasks that need to be tackled; while each task is assigned to a team of agents (denoted as \mathcal{A}) who collectively work towards the task. There is a plethora of team formation algorithms, referred to as TFA, that solve the team formation problem. A TFA takes as input data regarding agents' characteristics along with data regarding tasks' descriptions, and outputs a team-to-task allocation (denoted as g), i.e., a mapping from tasks to teams of agents. The team-maker is the one who invokes the TFA to generate an allocation. Throughout the paper we will be using as a running example the "classroom scenario": a teacher needs to split their students into teams that work on different projects each. As such, we will be considering algorithms as the one proposed in [14, 12], since it applies best in our classroom scenario.

In the context of team formation, Georgara et al. [13] propose a general scheme that demonstrates explanations using a many teams to many tasks TFA. Specifically, [13] proposes a general explanatory algorithm that *wraps* existing team formation algorithms in order to build *contrastive explanations* regarding a team-to-tasks allocation. As highlighted in [22], contrastive explanations are based on findings in the philosophical and cognitive sciences literature indicating that people are not interested in the causes leading to a particular outcome (in our case an allocation) per se, but, on the contrary, they are interested in the causes that explain a non-occurring outcome. In other words, people are interested in (and also tend to give) explanations regarding questions of the type "Why X instead of Y?". As such, a contrastive explanation provides the reasons why outcome X is preferred to another outcome Y.

A contrastive explanation within team-to-tasks allocation problems corresponds to information coming from the comparison between two allocations that justify why one is preferred to the other. In [13], the authors build explanations of the form: "If team A was assigned to task τ instead of B , it would result in task τ being assigned to a team (A) that is worse than its current team (B) with respect to property f ". According to [13] the TFA at hand forms teams based on some *desired properties*, while there is a way to measure the matching quality of a team being assigned to a task toward some desired property. Therefore, they exploit these desired properties in order to justify why one task assignment $\langle \text{team}_1, \text{task}_1 \rangle$ in an allocation is better than a task assignment $\langle \text{team}_2, \text{task}_2 \rangle$ in an alternative allocation. Now, in our classroom example, we consider the following four desired properties: (1) a team shall be *skilled* for its assigned task, (2) a team shall be diverse in terms of individuals' personalities, (3) a team shall be satisfied with their assigned task, and (4) a team shall be socially coherent. As such, here we use *individuals' features* (skills, personality) and *individuals' preferences* (over projects, over potential team-mates) in order to measure the matching quality of a team with a task *wrt.* each one of the desired properties.

With this desired properties in mind, the teacher uses the TFA which forms a teams-to-tasks allocation. Then a student, namely Beth, challenges the explana-




Fig. 1. General Framework for Privacy-Aware Explanations for Team Formation

tory algorithm (e.g., the one in [13]) with the following query: “Why is Jack in my team instead of Alex?”. Given the query, the EA (according to [13]) computes an alternative allocation that *enforces* Beth and Alex to work together. Then the EA computes the differences between the teams in both allocations for Beth. That is, let according to the current allocation Beth be working with Jack on project Maths Game, while according to the alternative allocation Beth is working with Alex on project Creative Writing. The EA would compare $\langle \text{Beth and Jack, Maths Game} \rangle$ against $\langle \text{Beth and Alex, Creative Writing} \rangle$ with respect to the properties: each team being (a) skilled for their assigned project; (b) diverse in terms of personality; (c) satisfied with their assigned project in terms of individuals’ preferences over projects; and (d) socially coherent in terms of individuals’ preferences over team-mates. Then the explanation that the EA builds would be: “If Alex was on your team instead of Jack, then you would be in a less diverse team in terms of personality than the the team you are currently in”. Notably, according to the EA, the desired property that justifies best why Beth should be working with Jack and not with Alex, is that of personality.

3 A General Framework for Privacy-Aware Explainable Team Formation

In this section we describe a general framework that combines team formation solutions, explanations over these solutions, and a mechanism for checking whether some explanation may cause a privacy breach. Assume we have a team formation scenario and a set of agents along with a set of task in our disposal. Moreover, let o be the ‘orchestrator’ or team-maker, i.e., the person who requests the forming of a team-to-tasks allocation using some team formation algorithm. A user is someone that challenges the teams-to-tasks allocation, and is either the team-maker or an agent. In this work we assume that each agent holds a view of the world which consists of: (i) *known facts* such as their own private information, the description of the tasks, and the teams-to-tasks allocation; and (ii) *beliefs* over other agents’ private information. Similarly, the team-maker also holds their own view of the world, consisting of some known facts and their beliefs over the agents’ private information. Figure 3 illustrates our proposed framework, which in a nutshell consists of the following components:

1. A team formation algorithm (TFA) that forms a teams-to-tasks allocation.
2. An explanatory algorithm (EA)—interacting with the TFA—that generates explanations regarding a teams-to-tasks allocation.
3. A privacy breach detector (PBD) that assesses whether an explanation may incur in privacy breaches. The PBD is composed of:
 - (a) a belief updater (BU) that computes posterior beliefs that the user is expected to form upon receiving an explanation; and
 - (b) a privacy checker (PC) that assesses whether the user’s expected posterior beliefs exceed a belief threshold.

In more details, the team-maker uses the TFA to solve a team formation problem and form an allocation; while the TFA notifies the team-maker and the agents with the allocation formed. As we mentioned in Section 2, there is a plethora of TFAs solving different team formation problems, as such depending on the problem at hand, one shall use the corresponding TFA. For example, a teacher in a classroom acts as the team-maker and uses a TFA to group their students (who correspond to agents) into teams in order to work on their mid-term projects (which correspond to tasks). The TFA computes the teams along with their allocation to projects. Thereafter, the TFA communicates the resulting teams and allocations to both the teacher and the students.

Then, say that some user challenges the TFA’s result. That is, a user may argue that there is a better allocation than the one yielded by the TFA. Hence, the user poses a question to the explanatory algorithm. For example, student Beth asks why Jack is in her team instead of her friend Alex. The EA processes the user’s question and, by interacting with the TFA, builds an appropriate explanation. For example, the EA builds the following explanation: “If Alex was on your team instead of Jack, then you would be in a less diverse team in terms of personality than the the team you are currently in”.

Next, the EA passes the generated explanation to the privacy breach detector, and in particular to the belief updater. As mentioned before, each agent holds knowledge regarding the world, and beliefs over other agents’ private information. The BU is responsible for exploiting the information conveyed by an explanation, combining it with the user’s knowledge and current beliefs in order to extract valuable conclusions. Specifically, the BU follows a theory of mind [11] on the user to simulate the reasoning that the user is expected to follow (based on the user’s knowledge and beliefs). As a result, the BU forms an updated version of beliefs which the user is expected to reach after receiving the explanation. For example, Beth is expected to update her beliefs on Alex’s and Jack’s personalities based on the explanation from the EA.

After that, the BU passes the expected posterior beliefs to the privacy checker. The PC is responsible for assessing whether the user’s expected posterior beliefs exceeds the belief threshold ε . The belief threshold corresponds to a maximum probability with which a user may believe that some agent’s information is true, without violating this agent’s privacy. For example, with a belief threshold $\varepsilon = 0.5$, if Beth is expected to update her beliefs that Alex is of personality role ‘leader’ to 0.3, then this *causes no* violation of Alex’s privacy. On the other hand,

if Beth is expected to update her beliefs that Jack is of personality ‘implementer’ to 0.7, then this *causes* a violation of Jack’s privacy.

Finally, the privacy checker outputs an answer for the explanatory algorithm. Specifically, the PC responds with an appropriate message indicating whether the explanation is *safe* if our PBD detected no privacy breaches on private information, or otherwise. Depending on the PC’s response, the explanatory algorithm either provides the explanation to the user, or *handles* this situation by e.g. computing a different explanation or denying to answer due to a privacy breach.

4 Representing knowledge and beliefs

In this section we discuss how to represent *knowledge* and *beliefs* used within our framework (see Figure 3). Recall that, as mentioned in Section 3, both the agents and the team-maker hold a view of the team formation problem which consists of known facts and beliefs. *Knowledge* corresponds to known facts that an agent has over the team formation scenario. Such known facts include the tasks’ description and the team-to-tasks allocation published by the TFA. Moreover, for an agent, known facts also include their own personal characteristics, i.e., this agent’s features and preferences. Besides knowledge, an individual can form over others. Specifically, an individual forms beliefs regarding knowledge they do not own, i.e., beliefs over another agent’s personal characteristics.

Agent’s Knowledge. An agent holds *knowledge* that can be either *private* or *public*. Given an agent $a \in \mathcal{A}$, their *private knowledge* refers to characteristics that comprise agent’s a own profile i.e., their features and preferences. Each agent holds their own private knowledge, withheld from anyone else. For example, “John is capable in Maths” corresponds to agent John’s private knowledge. *Public knowledge* refers to the tasks made public by the team-maker and the team-to-tasks allocations received from the TFA. All agents at the outset share the same public knowledge, i.e. public knowledge is common to all agents. For example, “John has been allocated to work on the Maths-Game task” is an example of public knowledge.

We represent knowledge using *first-order predicates* with *ground terms*. For an agent $a \in \mathcal{A}$ we denote the private knowledge of a as Γ_a corresponding to a set of first-order predicates with ground terms referring to a . For instance, predicate $acquire(John, maths) \in \Gamma_{John}$ corresponds to some of John’s private knowledge. In our running example, the predicates that exist in agent’s a private knowledge are: $acquires(a, skill)$, $personality(a, role)$, $wants_to_work_on(a, \tau)$, and $wants_to_work_together(a, b)$, where $skill$ is some skill that agent a acquires, $role$ is agent’s a personality role, τ is some task that a wants to work on, and $b \in \mathcal{A}$ is some agent that a wants to work with. We denote with Γ_τ the public knowledge that each agent initially holds regarding task τ —for example, $size(Maths - Game, 3) \in \Gamma_{Maths - Game}$; and with Γ_g the public knowledge that each agent initially holds regarding the team-to-tasks allocation—for example, $worksOn(g, John, \tau) \in \Gamma_g$, which is read as “According to allocation g , John is assigned to work on the Maths-Game task”.

Team-Maker’s Knowledge The team maker only holds public knowledge regarding the tasks’ description and team-to-tasks allocations. Thus, the team-maker’s knowledge is $\Gamma_o \equiv \Gamma_\tau \cup \Gamma_g$.

Agents’ Beliefs. Each agent holds beliefs over other agents’ private knowledge. That is, an agent sets a probability with which they believe that some private knowledge of another agent is true. For example, let Beth believe that John is knowledgeable in maths with probability 0.7, this comprises Beth’s belief over some of John’s private knowledge. This belief is in fact a probability over a predicate in Γ_{John} , i.e., $P[acquires(John, maths)] = 0.7$. Thus, an agent’s a beliefs correspond to a probability function over predicates in $\bigcup_{b \in \mathcal{A}} \Gamma_b$.

Team-maker’s Beliefs. The team-maker holds beliefs over the agents’ private information as well. Similarly, the team-maker’s beliefs correspond to a probability function over predicates in $\bigcup_{a \in \mathcal{A}} \Gamma_a$.

5 Inference Rules

Here we discuss about the *inference rules* used by our model within the privacy breach detector (see Figure 3). We use ‘IF-THEN’ rules that guide the BU component to reason over new information deriving from an explanation. Specifically, we discern rules that (i) determine when a team satisfies a desired property, and (ii) interpret a comparison described in an explanation.

Considering our classroom example, we have one rule per desired property to determine when a team satisfies this property. For example, such a rule is: “IF the team members are of different personality roles THEN the team is diverse”, which using first order predicates is written as: $\forall x, y \forall p \text{ inTeam}(x, A) \wedge \text{inTeam}(y, A) \wedge \text{personality}(x, p) \wedge \neg \text{personality}(y, p) \Rightarrow \text{isDiverse}(A)$. We also have rules for interpreting a comparison described in explanations. The comparison of an explanation is in the form of “team A assigned to τ satisfies property f , while team B assigned to σ does not” or “both team A assigned to τ and team B assigned to σ (do not) satisfy property f ”; which using first-order predicates is written as: $\text{isAssigned}(A, \tau) \wedge \text{isAssigned}(B, \sigma) \wedge \text{isBetter}(A, B, f) \Rightarrow \text{satisfies}(A, \tau, f) \wedge \neg \text{satisfies}(B, \sigma, f)$ and $\text{isAssigned}(A, \tau) \wedge \text{isAssigned}(B, \sigma) \wedge \text{isEqual}(A, B, f) \Rightarrow (\text{satisfies}(A, \tau, f) \wedge \text{satisfies}(B, \sigma, f)) \vee (\neg \text{satisfies}(A, \tau, f) \wedge \neg \text{satisfies}(B, \sigma, f))$.

Given these rules, we can handle the process of inference with a rule-based forward reasoner [27]; while the inference is used to update the beliefs that the explaineo holds over private information of the agents appearing in the explanation, following a theory of mind approach [11].

6 Conclusions

In this paper we tackled the challenge of preserving privacy when providing explanations within the multi-agent setting of team formation. We argue that providing explanations should guarantee that agents’ private information is not

disclosed. Towards this, we propose a general framework that combines team formation solutions and explanations over these solutions, while it detects potential privacy breaches upon offering explanations. In particular, we put forward a privacy breach detector that complements an explanatory algorithm as the one proposed in [13], and assesses the explanations built *wrt.* privacy breaches.

References

1. Adadi, A., Berrada, M.: Peeking inside the black-box: A survey on explainable artificial intelligence (xai). *IEEE Access* **6**, 52138–52160 (2018)
2. Anagnostopoulos, A., Becchetti, L., Castillo, C., Gionis, A., Leonardi, S.: Power in unity: Forming teams in large-scale community systems. pp. 599–608 (01 2010)
3. Andrejczuk, E., Berger, R., Rodríguez-Aguilar, J.A., Sierra, C., Marín-Puchades, V.: The composition and formation of effective teams: computer science meets organizational psychology. *Knowledge Eng. Review* **33**, e17 (2018)
4. Antognini, D., Musat, C., Faltings, B.: Interacting with explanations through critiquing. In: Zhou, Z.H. (ed.) *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*. pp. 515–521 (8 2021), main Track
5. Boixel, A., Endriss, U.: Automated justification of collective decisions via constraint solving. In: *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. p. 168176. *AAMAS '20* (2020)
6. Borg, A., Bex, F.: Contrastive explanations for argumentation-based conclusions. In: *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems* (2022)
7. Capezzuto, L., Tarapore, D., Ramchurn, S.D.: Anytime and efficient coalition formation with spatial and temporal constraints (2020)
8. Carvalho, D.V., Pereira, E.M., Cardoso, J.S.: Machine learning interpretability: A survey on methods and metrics. *Electronics* (2019)
9. Crawford, C., Rahaman, Z., Sen, S.: Evaluating the efficiency of robust team formation algorithms. In: Osman, N., Sierra, C. (eds.) *Autonomous Agents and Multiagent Systems*. pp. 14–29. Springer International Publishing, Cham (2016)
10. Doilovi, F.K., Bri, M., Hlupí, N.: Explainable artificial intelligence: A survey. In: *2018 41st International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*. pp. 0210–0215 (2018)
11. Frith, C., Frith, U.: Theory of mind. *Current biology* **15**(17), R644–R645 (2005)
12. Georgara, A., Rodríguez-Aguilar, J.A., Sierra, C.: Allocating teams to tasks: An anytime heuristic competence-based approach. In: Baumeister, D., Rothe, J. (eds.) *Multi-Agent Systems - 19th European Conference, EUMAS 2019, Düsseldorf, Germany, September 14-16, 2022*
13. Georgara, A., Rodríguez-Aguilar, J.A., Sierra, C.: Building contrastive explanations for multi-agent team formation. In: *Proceedings of the 21st International Conference on Autonomous Agents and MultiAgent Systems. AAMAS '22* (2022)
14. Georgara, A., Rodríguez-Aguilar, J.A., Sierra, C., Minch, O., Kazhemiakin, R., Palmero Approsio, A., Pazzaglia, J.C.: An anytime heuristic algorithm for allocating many teams to many tasks. In: *Proceedings of the 21st International Conference on Autonomous Agents and MultiAgent Systems* (2022)
15. Goodman, B., Flaxman, S.: European union regulations on algorithmic decision-making and a right to explanation. *AI Magazine* **38**(3), 50–57 (Oct 2017)

16. Holzinger, A.: From machine learning to explainable ai. In: 2018 World Symposium on Digital Intelligence for Systems and Machines (DISA) (2018)
17. Kleinerman, A., Rosenfeld, A., Kraus, S.: Providing explanations for recommendations in reciprocal environments. In: Proceedings of the 12th ACM Conference on Recommender Systems. p. 2230. RecSys '18, Association for Computing Machinery, New York, NY, USA (2018)
18. Kökciyan, N., Yolum, P.: P r i g uard: A semantic approach to detect privacy violations in online social networks. *IEEE Transactions on Knowledge and Data Engineering* **28**(10), 2724–2737 (2016)
19. Kraus, S., Azaria, A., Fiosina, J., Greve, M., Hazon, N., Kolbe, L., Lembcke, T.B., Muller, J., Schleibaum, S., Vollrath, M.: Ai for explaining decisions in multi-agent environments. *Proceedings of the AAAI Conference on Artificial Intelligence* (2020)
20. Kunkel, J., Donkers, T., Michael, L., Barbu, C.M., Ziegler, J.: Let me explain: Impact of personal and impersonal explanations on trust in recommender systems. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA (2019)
21. Lipton, Z.C.: The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue* **16**(3), 3157 (2018)
22. Miller, T.: Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* **267**, 1–38 (2019)
23. Mohseni, S., Zarei, N., Ragan, E.D.: A multidisciplinary survey and framework for design and evaluation of explainable ai systems (2020)
24. Mosca, F., Such, J.M.: ELVIRA: An Explainable Agent for Value and Utility-Driven Multiuser Privacy, p. 916924. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2021)
25. Nardi, O., Boixel, A., Endriss, U.: A graph-based algorithm for the automated justification of collective decisions. In: Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems. AAMAS '22 (2022)
26. Puiu, A., Vizitiu, A., Nita, C., Itu, L., Sharma, P., Comaniciu, D.: Privacy-preserving and explainable AI for cardiovascular imaging. *Studies in Informatics and Control, ISSN 1220-1766* **30**(2), 21–32 (2021)
27. Rattanasawad, T., Saikaew, K.R., Buranarach, M., Supnithi, T.: A review and comparison of rule languages and rule-based inference engines for the semantic web. In: 2013 International Computer Science and Engineering Conference (ICSEC)
28. Rosenfeld, A., Richardson, A.: Explainability in human-agent systems. *Autonomous Agents and Multi-Agent Systems* **33**(6), 673–705 (Nov 2019)
29. Samek, W., Müller, K.R.: *Towards Explainable Artificial Intelligence*, pp. 5–22. Springer International Publishing, Cham (2019)
30. Sörries, P., Müller-Birn, C., Glinka, K., Boenisch, F., Margraf, M., Sayegh-Jodehl, S., Rose, M.: Privacy needs reflection: Conceptual design rationales for privacy-preserving explanation user interfaces. *Mensch und Computer 2021-Workshopband*
31. Sovrano, F., Vitali, F., Palmirani, M.: Making things explainable vs explaining: Requirements and challenges under the gdpr. In: *AI Approaches to the Complexity of Legal Systems XI-XII: AICOL International Workshops 2018 and 2020: AICOL-XI@JURIX 2018, AICOL-XII@JURIX 2020, XAILA@JURIX 2020, Revised Selected Papers*. p. 169182. Springer-Verlag, Berlin, Heidelberg (2020)
32. Such, J.M., Criado, N.: Multiparty privacy in social media. *Communications of the ACM* **61**(8), 74–81 (2018)
33. Zhang, Y., Chen, X.: Explainable recommendation: A survey and new perspectives. *Foundations and Trends in Information Retrieval* **14**(1), 1–101 (2020)