


Article

User-Centric Privacy for Identity Federations Based on a Recommendation System

Carlos Villarán [†] and Marta Beltrán ^{*,†} 

Department of Computing, ETSII, Universidad Rey Juan Carlos, c/Tulipan s/n, 28933 Mostoles, Spain; carlos.villaran@urjc.es

* Correspondence: marta.beltran@urjc.es

† These authors contributed equally to this work.

Abstract: Specifications such as SAML, OAuth, OpenID Connect and Mobile Connect are essential for solving identification, authentication and authorisation in contexts such as mobile apps, social networks, e-commerce, cloud computing or the Internet of Things. However, end-users relying on identity providers to access resources, applications or services lose control over the Personally Identifiable Information (PII) they share with the different providers composing identity federations. This work proposes a user-centric approach based on a recommendation system to support users in making privacy decisions such as selecting service providers or choosing their privacy settings. The proposed Privacy Advisor gives end-users privacy protection by providing personalised recommendations without compromising the identity federations' functionalities or requiring any changes in their underlying specifications. A proof of concept of the proposed recommendation system is presented to validate and evaluate its utility and feasibility.

Keywords: identity infrastructures; federated identity management; privacy; recommendation system



Citation: Villarán, C.; Beltrán, M. User-Centric Privacy for Identity Federations Based on a Recommendation System. *Electronics* **2022**, *11*, 1238. <https://doi.org/10.3390/electronics11081238>

Academic Editors: Diana Berbecaru, Detlef Hühnlein and Costin Badica

Received: 9 March 2022

Accepted: 8 April 2022

Published: 14 April 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Federated Identity Management (FIM) specifications allow resource, application or services providers (Relying parties or RPs) to solve authentication or authorisation of end-users trusting in the authentication performed by an external Identity Provider (IdP).

Users of federated identity management are comfortable with these mechanisms because they avoid creating a local account in each resource, application or service. It is only necessary to have one in a few identity providers. However, they make privacy-related decisions every time they enrol at a new provider, create an account, choose their privacy settings, or use this provider when accessing an online resource, service or application [1]. As a result, privacy decision making is a significant burden for end-users who usually rely on the default configuration and settings [2]. It must be considered that they might not be the most appropriate since IdPs are, on many occasions, large technology companies such as Facebook or Google (providers of social login based on OAuth [3], or OpenID Connect [4]) or mobile network operators (providers of identity management services based on Mobile Connect [5]) with their own interests and business models.

This work relies on end-user engagement in their own privacy protection: an appropriate communication of privacy risks in a given scenario can prevent privacy threats or mitigate their impacts [6]. A well-designed recommendation system can bring about this engagement [7], a Privacy Advisor capable of informing end-users about data protection practices within identity federations, providing personalised advice, and even acting on behalf of the user in specific cases. This approach extends traditional privacy architectures beyond the RP and the IdP; end-users are also involved in their privacy protection. The personalisation of recommendations is essential because previous research has found that perceived personalisation significantly increases users' intentions to follow

provided guidance by increasing their cognitive and emotional trust in recommendation systems [8–10].

The Privacy Advisor can be offered by the identity providers themselves as a value-added service or, if their neutrality is not trusted, can be a new agent in the federations, utterly independent of the three traditional roles; for example, groups protecting the rights of Internet users or governmental agencies are offered by non-profit organisations.

This paper's main contributions are: (1) a model that adds a new role to identity federations, the Privacy Advisor (PAdv), a recommendation system devoted to giving end-users actionable and understandable privacy recommendations and supporting them in their decisions; (2) a set of this recommendation system's required capabilities; (3) a modular architecture that is easy to implement, modify and extend to provide these capabilities without modifying current federated identity management flows; and (4) a proof of concept of the PAdv to provide a first validation and evaluation of the proposed model.

The rest of this paper is organised as follows. Section 2 summarises previous work on federated identity management, privacy and user-centric privacy protection. This discussion of previous work leads to the motivation of our proposal in Section 3. This section also introduces the proposed model, identifies the required capabilities of the Privacy Advisor, and shows its proposed architecture. Section 4 presents the design of the Privacy Advisor and all the modules that are part of it. Section 5 explains how an advisor's first prototype has been implemented and discusses the proposed model's first validation and evaluation. Finally, Section 6 provides conclusions and some interesting lines for future research.

2. Related Work

2.1. On Identity Federations and Privacy

Different researches have shown in the past how the utilisation of federated identity management poses privacy threats for users [11,12]. However, this paradigm benefits both end-users and application or service providers. There has been a significant research effort focused on proposing different privacy-preserving mechanisms to mitigate these privacy threats.

In [13,14], the user is able to control her PII by using personalised policies. A Privacy Controller or PC allows users in [13] to check and modify PII and also to control its processing and sharing. This mechanism gives users the ability to know how the PII disclosure is performed. In [14], the proposed solution is very similar but adapted to mobile environments by using a new element, the Mobile Information Service Broker. The work presented in [15] covers the improvement of privacy in the SAML 2.0 authentication process, defining a Privacy Engine Module responsible for monitoring user data usage. In [16], a selective PII disclosure is proposed, using brokers and a Proxy Re-Encryption. The brokers manage the IdP and RP centrally, and the Proxy Re-Encryption allows changing the encryption key, over a ciphered element, without ever decrypting that element.

Finally, PRIMA [17] proposes a different approach, an authentication flow that does not require any interaction between RPs and identity providers, avoiding users profiling. Moreover, it enables controlled disclosure of users' PII.

As can be seen, the work that has so far attempted to prevent privacy threats arising in identity federations has focused either on preventing the sharing of personal data with providers (or doing so with encrypted data) or ensuring transparency and control of this sharing (personalised sharing policies, real-time monitoring, selective disclosure). Adding a new agent to the federations, an intermediary such as a privacy controller, broker, or proxy, to add these capabilities is exciting and is repeated in many proposals. However, no work allows the informed self-management of this privacy by putting the user at the centre of his or her decisions, nor that takes advantage of this new agent beyond the capabilities mentioned above.

2.2. On User-Centric Privacy Protection

Different authors have proposed recommendation systems in the form of privacy assistants or advisers to help users find their appropriate privacy settings. Previous research has demonstrated that recommendation systems have a significant effect on users' behaviour and that users find short lists of recommended actions helpful [18,19]. It could be considered using the new agent mentioned in the previous section as such a recommender. However, it is something that has not yet been tested in identity federations to the best of our knowledge.

However, Table 1 shows a summary of different user-centric recommendation systems for privacy performed to compare the proposed solution with previous works in other application domains. Due to the user-centric approach, all the analysed works have a user-tailored output. All the analysed recommendation systems have been categorised in this table considering their application domain; additional columns summarise the essential characteristics of the recommendation systems. The "User customisation" column refers to the user's ability to select or modify some of the recommendation systems' evaluation criteria. The "Ease of adoption" refers to the volume or complexity of actions the user has to perform to use the recommendation system. The "User-friendly" column expresses the ease to use and understand the system output, i.e., the provided recommendation. The "Ease to integrate" column denotes if the recommendation system is easy to integrate within the application domain without substantial changes in the underlying specifications, protocols or implementations.

Table 1. User-centric recommendation systems for privacy: comparison.

Ref.	Domain	User Cust.	Ease of Adopt.	User-Friendly	Ease to Int.
[20]	Social net.		X		X
[21]	Social net.		X	X	X
[22]	Social net.		X	X	X
[23]	Social net.	X	X	X	X
[24]	Social net.	X	X		
[25]	Web-based		X	X	
[26]	Web-based	X	X		
[27]	Web-based		X		
[28]	Mobile apps	X			X
[29]	Mobile apps	X			X
[30]	Mobile apps		X		
[31]	Mobile apps	X			X
[32]	IoT	X	X		
Our Work	FIM	X	X	X	X

It can be seen that the solutions proposed in the literature are an excellent starting point for the contributions intended to be made in this research. However, they often suffer from a lack of customisation, a lack of user-friendliness and, above all, a lack of ease of integration with underlying protocols, specifications and technologies. This last aspect is crucial, as it makes it very difficult for the proposals made to be transferred to real production environments. It is unrealistic to assume that different providers will adapt their infrastructures and deployments to match what the proposed solutions require.

It should also be noted that the works analysed are proposed in application domains other than identity federations, such as social networks, web, mobile or IoT. There are two key aspects that differentiate our research from these previous works:

1. Previous work focuses on making recommendations about a provider, server, app or device. In our case, we have to consider multi-actor contexts where the recommendations made consider the RP and the IdP.
2. Previous work focuses on making recommendations only once, at the beginning. For example, when the user signs up for a service or installs an app for the first time. In our case, recommendations are continuous and real-time, each time the user uses the IdP to authenticate to a resource, application or service.

2.3. On Recommendation Systems

A recommendation system is a solution that assists an individual who lacks knowledge or experience about a specific topic when making a choice or a decision [33]. This assistance is often based on the user's interests or preferences, the possible alternatives, the relations between them, or other individuals' behaviours. This kind of system has been the basis of an active research line for more than 20 years.

There are a plethora of application domains for recommendation systems. One has already been analysed in the previous section, recommendations in the area of privacy. They can also be used in the area of cybersecurity [34]. However, perhaps the most common applications are in the domains of Streaming Services, Social Network Services, Tourism Services, E-Commerce Services, Healthcare Services, Education Services and Academic Information Services [35].

Regarding the underlying techniques and models used for building recommendation systems, they are usually categorised as Collaborative Filtering, Content-Based, Knowledge-Based and Hybrid.

Recommenders based on Collaborative Filtering recommend items to individuals based on similar individuals' opinions and last choices (nearest neighbours, with similar preferences or tastes) [36]. These techniques can be applied in domains where there is not much information about the individual or the item or where information regarding sentiments or opinions (helpful to build an individual's profile) may be challenging to extract automatically. On the other hand, these techniques are susceptible to the similarity measure used to quantify the degree of similarity between individuals [37] and may have issues when one individual belongs to more than one group.

Recommenders based on Content-Based Filtering recommend items to individuals with attributes similar to those that users liked in the past and recommend them based on the information of the items [38]. These techniques can be easily adapted when the preference of individuals change and are easily explainable. However, they tend to be "overspecialised" (not able to suggest items different from the items chosen before); and require data about the individual's profile (past choices, opinions) and the items to make accurate predictions.

Knowledge-based recommenders do not base their recommendations on traditional data about individuals and items. Their recommendations are based on explicit rules about the applications' domain and about the context [39]. Therefore, they require a deep knowledge of the application domain to be built. On the other hand, they perform well in complex domains, are reliable, and do not suffer traditional problems such as the cold start (when the system cannot infer anything about individuals or items if there is not enough data yet).

Finally, Hybrid recommenders combine features of the three aforementioned categories, trying to take advantage of the best of each of them [40].

3. The Proposed Model

Since recommendation systems have been demonstrated to be a promising approach to support users when making privacy decisions in other application domains, this work proposes an approach complementary to the research performed in the past. While providers work to incorporate privacy-preserving mechanisms, users must have mechanisms in place that allow them to make the most appropriate decisions to protect their privacy. Further-

more, these mechanisms must be global, capable of working with different specifications while personalised, considering different user profiles/categories and preferences. Moreover, they must be highly customised, easy to adopt and user-friendly. Furthermore, easy to integrate with all possible identity management flows while considering the different aspects of privacy protection proposed in previous research: design patterns and compliance analysis, collaboration between users and their feedback, and privacy scoring or reputation systems.

3.1. Motivation and Overview

The Privacy Advisor's primary purpose is to show users how their privacy is affected when using a service provider (an RP) within an identity federation. The PAdv can make recommendations to help users decide on whether to use a service or not. Alternatively, use it with certain restrictions or only after changing account configuration at the IdP. These recommendations are generated, providing users high control over the inputs requested (privacy profile), its processes (data collected from RPs to generate recommendations) and output (provided recommendations and their meanings).

It is relevant to propose such a tool within identity federations because most exchanges of sensitive information about the end-user take place without the user's knowledge, directly between the IdP and the RP. When the RP manages identities natively, interacting directly with the end-user, the user is usually more aware of the type of data it shares with the RP.

Considering the background provided in Section 2.3 and the motivation of this work, the PAdv relies on content-based filtering mechanisms, relying on a set of discrete and tagged features of users and RPs to provide a recommendation. This advisor has the following capabilities:

1. To collect a minimum set of end-users privacy features regarding PII protection.
2. To gather updated information about how a resource, application or service provider handles the users' PII. The PAdv can use different data sources to get all the information needed.
3. To provide a recommendation about privacy protection.
4. To show all the information in a friendly and illustrative way to the users, making it easier to interpret the PAdv recommendations to have a tangible impact on users' decisions.
5. To offer more detailed information about the recommendation results in layers for advanced or more aware users. The PAdv should be able to work with different granularity levels or detail levels when making recommendations.

These capabilities allow us to increase transparency and control within identity federations, two essential pillars of privacy often forgotten in this context.

3.2. High Level Design

This work proposes an entirely modular architecture for the Privacy Advisor, described in Figure 1, to provide the capabilities mentioned above. Each module has its own functionality and may run in a distributed manner: the modules can belong to different organisations or companies if desirable. For example, relying on REST, GraphQL or JSON APIs. As mentioned before, the proposed user-centric solution could be offered by an identity provider or other organisations. Alternatively, they could all collaborate in supporting the user in her privacy decision-making.

The PAdv could be implemented in different ways, for example as a plugin, service or app. The advantages of each alternative do not depend on the use cases, but on who is in charge of this recommender system and for what type of user. For example, if it is offered by a mobile network operator associated with Mobile Connect, it probably makes sense to offer it as a mobile app, whereas if a government agency offers it for social login users, it probably makes more sense to implement it as a service or plugin.

Integrating new modules is also easy because they are designed to avoid strong dependencies. As shown in Figure 1, there is a Feature Collection module responsible for collecting and storing end-users profiles (step 1). The user's features help the Privacy Advisor to adapt the recommendations to the users' needs, devices, etc. When the end-user interacts with an RP (step 2), the Recommendation module receives a request and asks for the Data Collection module the required information to produce this recommendation, consulting the end-user profile stored in the Feature Collection module to personalise it. Once it has received all the information, it returns a response to the user with a recommendation (step 3).

The recommendations, in the first layer, are presented with a traffic light colour code: green (OK, Good), yellow (partially OK, Caution) and red (NO OK, Bad), correlating the information gathered from the Data Collection module and the user's profile. There is also an Unknown category with a grey colour when it is impossible to provide a personalised or well-grounded recommendation due to lack of information, user features or RP data.

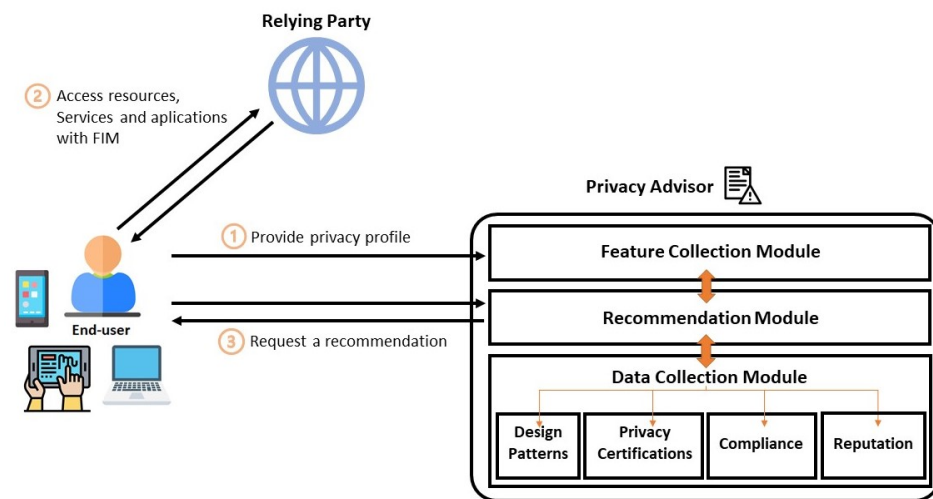


Figure 1. Privacy Advisor architecture.

3.3. Integration with Identity Federations

The addition of this new agent, the Privacy Advisor, does not affect the authentication or authorisation flows defined by the federated specifications, shown in Figure 2 from step 3. This aspect is essential if we want to use it in production environments without modifying these specifications or the products that implement them. It should be noted that the RP is the Relying Party (the provider of a resource, application or service), the IdP is the Identity Provider (where the user has an account), and the End-User is the individual asking for access to the resource, application or service and receiving a recommendation from the PAdv.

Two different flows are triggered when a user requests access to an RP and needs to authenticate using a federated scheme. The first is the usual, the authentication or authorisation flow (steps 3 to 8 in Figure 2). The second is a request to the PAdv asking for a recommendation (step 1 in Figure 2). The PAdv responds to end-user requests with a tailored recommendation, including the comparison between the privacy results from this service and the user's privacy profile (step 2 in Figure 2). These two steps can be carried out before or during the authentication process (as mentioned before, step 3 and further in Figure 2). The example shown in Figure 2 uses the Authorisation Code flow of OpenID Connect. However, it could use any flow from the OpenID Connect specification or similar federated specifications (SAML, OAuth, Mobile Connect, etc.).

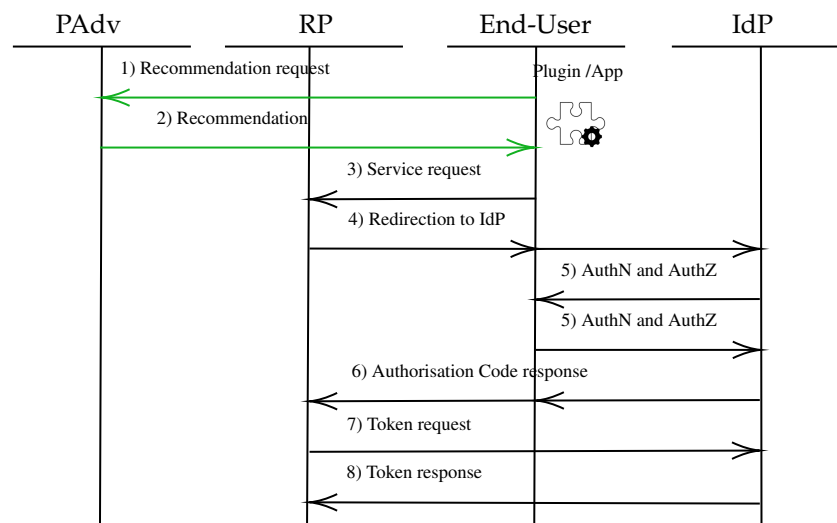


Figure 2. Example of an authentication flow with the PAdv and OpenID Connect.

The PAdv can operate in two different ways depending on the user preferences. The first one, as shown in Figure 2, triggers the authentication process once the user has revised and accepted the privacy recommendation. This way to proceed may be intrusive, but it guarantees that the user knows how accessing a particular RP affects her privacy. The second one runs parallel with the authentication process (for example, in a new tab or window in the browser). This process is entirely independent, but it does not force the user to check the recommendation before interacting with the requested resource, application or service.

4. Privacy Advisor Modules Design

4.1. Feature Collection Module

The collection of features and attributes is a guided process that takes around 20/30 min the first time and can be updated as many times as the end-user needs. This process allows the PAdv to build a user's profile explicitly asking:

- Which identity providers the user is working with (from a list).
- Data shared with each of them, specifically PII, and kind of account.
- Devices from which they are used.
- Categories of accessed service providers (RPs) and the most frequently used or visited set. A simple classification algorithm is used to classify RPs using the SimilarWeb categories [41]: Adult, Arts and Entertainment, Computers Electronics and Technology, News and Media, etc.

All these attributes allow the PAdv to model the exposure level of a specific end-user. Four discrete levels are used and summarised in Table 2: improbable (1), low (2), high (3), or very high (4). Thus, a user who, for example, uses a single IdP with any known security vulnerabilities, to whom she has only provided her contact details, and who uses a laptop to access low-sensitive categories of RPs, will be labelled with a low level of risk or exposure. On the contrary, a user who uses several IdPs, some with known vulnerabilities, has provided credit card details to some of them, regularly uses mobile devices and accesses numerous sensitive RPs (for example, finance, e-commerce or health), will be labelled with a very high level of risk or exposure.

Table 2. Exposure level.

Score	Explanation
1—Improbable	Using no IdP with known vulnerabilities (not solved) and sharing no sensitive data with IdPs: name, surname, picture, personal identification number, etc.
2—Low	Using no IdP with known vulnerabilities (not solved) and sharing sensitive data with IdPs: phone number, email, credit card, bank account, etc.
3—High	Using at least one IdP with known vulnerabilities (not solved) and sharing no sensitive data with IdPs: name, surname, picture, personal identification number, etc.
4—Very high	Using at least one IdP with known vulnerabilities (not solved) and sharing sensitive data with IdPs: phone number, email, credit card, bank account, etc.

In addition, the user shows, through examples, which service providers (RPs) she would like to allow different levels of access to her PII. These levels allow the PAdv to model the end-user privacy awareness depending on her choices using four levels again: engaged (1), committed (2), interested (3), and unaware (4).

End-users are tagged with their Inherent risk score ($exposure \times awareness$), with values ranging from 1 ($improbable \times engaged$) to 16 ($very\ high \times unaware$). This risk does not change unless the user changes her profile, hence the term “inherent” risk. This approach based on a scoring system and a risk matrix has been selected because different standardisation bodies and risk management frameworks in the security and privacy fields have shown its benefits in the past. Including many devoted to privacy impact assessments and similar risk quantification procedures [42,43].

The inputs obtained from this process are summarised in a control panel and can be detailed and edited by navigating through different tabs and screens. A heat map, with the exposure level in one dimension and the end-user privacy awareness in the other, is used as a visualisation mechanism because it is concise and easy to understand. These maps are also valuable for comparison with other users, highlighting the most common behaviour of other peers and potential improvements to decrease the Inherent risk score. All these aspects are helpful to improve user awareness.

4.2. Data Collection Module

The Recommendation module accesses this module to produce its recommendation correlating the collected data and the end-user profile. The Data Collection module applies the classifier mentioned before to categorise the RP: Adult, Arts and Entertainment, Business and Consumer Services, Community and Society, Computers Electronics and Technology, E-commerce and Shopping, Finance, Food and Drink, Gambling, Games, Health, Heavy Industry and Engineering, Hobbies and Leisure, Home and Garden, Jobs and Career, Law and Government, Lifestyle, News and Media, Pets and Animals, Reference Materials, Science and Education, Sports, Travel and Tourism and Vehicles. A similar ranking could be carried out using other APIs such as the Alexa top site; the categories differ little from each other.

Furthermore, in this work, four information-gathering components are proposed for this Data Collection module: Design patterns revision, Privacy certifications, Compliance and Reputation.

These components have been designed to be independently modified or replaced. New components can be added too, without affecting the proposed model. Each component receives a different input from an RP, produces a result, and returns it to the Data Collection module. Different combinations of these components can be used depending on the category of the RP being assessed and its available information.

4.2.1. Design Patterns Revision

Dark patterns are based on ambiguous texts, opt-out options, distraction with pop-ups, colours and other interface attributes, hiding options or relevant information that organisations do not want the user to locate quickly (for example, to unsubscribe), etc. Therefore, these patterns are carefully designed to trick users into doing things that are not necessarily good or positive for them.

This component of the PAdv has to identify which dark patterns are used by an RP. With this information available in the Privacy Advisor, the users can be more aware of their interaction with this specific RP. Therefore, they can avoid the purpose of the dark patterns or even decide to use a more privacy-respectful alternative.

The proposed component should be able to identify at least the following dark pattern categories identified in [44]: Tricky Questions, Sneak into Basket, Roach Motel, Privacy Zuckering, Misdirection, Bait and Switch, Confirm shaming, Disguised Ads, Forced Continuity and Friend Spam.

4.2.2. Privacy Certifications

Privacy seals and certifications allow providers to certify the fulfilment of particular privacy and data protection requirements. These seals and certifications usually imply external audits at the beginning and periodically in the renewal process.

The Privacy Advisor gathers this valuable information for end-users when making decisions because a third party certifies privacy protection levels at a particular RP. This module's main functionality is to collect these privacy seals and certifications, list them, and inform the user about their implications if they are active.

This work proposes checking, at least, the following certifications: European Privacy Seal (EuroPriSe) [45], TrustArc company [46] certifications (APEC Cross Border Privacy Rules (CBPR), Enterprise Privacy & Data Governance Practices and Data Collection), ePrivacy [47] (ePrivacySeal and ePrivacyApp). These certifications have been selected because they are related to privacy in the EU and its relations with third parties; others could be easily added if required.

4.2.3. Compliance

The specifications used to solve federated identity management include specific sections devoted to security and privacy. The European GDPR, eIDAS and PSD2 regulations also include provisions that providers should comply with.

This component's main functionality is checking the fulfilment of the specifications' security and privacy best practices and the compliance with the mentioned regulations. When the Data Collection module asks for an RP evaluation, this component responds with an assessment specifying every compliance verification.

4.2.4. Reputation

This component is optional within the PAdv design but allows users to determine how reliable a service is given the opinion of other independent agents (other providers, users, Privacy Advisors, etc.). Therefore, this component quantifies the level of trust of a service by using third party evaluators (other components of the identity federation), understanding this trust as their confidence in the ability of the assessed provider of operating without threatening users' privacy.

This component requests a reputation assessment of a service to trusted third parties (RP, IdP, PAdv, etc.). These elements respond in a machine-readable format with a reputation score V_i between 0 and 100 based on the past behaviour of the assessed provider and any additional information about the reason for this assessment. The Privacy Advisor may not trust the evaluation of all the requested parties. This work proposes a weight measurement system with weight W_i (between 0 and 1) for each consulted party. The more weight a party has, the more reliable it is for a specific end-user. The reliance can be configured in

the user's preferences to give the users the chance to represent their trustworthiness in the parties integrated with this component.

The final reputation score is obtained with the following expression:

$$Reputation = \sum_{i=1}^n V_i \cdot W_i \quad (1)$$

Building these provider reputation systems and deciding the weights to be assigned to each evaluator's score is beyond the scope of this paper. However, it seems essential to include the possibility of taking into account this type of reputation system in the PAdv since in complex contexts involving a multitude of providers such as Cloud Computing [48] or IoT [49,50] (both similar to identity federations in terms of distribution and heterogeneity) it is a trend that is gaining more and more strength [51], even recurring to decentralised reputation systems based on blockchain [52].

4.3. Recommendation Module

End-users with different knowledge, preferences and perceptions of risk may be users of the recommendations provided by the PAdv.

The diversity and complexity of potential use cases in the considered domain make the goal of the proposed model, an informed end-user decision, challenging. The following list of characteristics is essential for the recommendation to have the desired effects:

- **Personalisation:** The provided recommendation must be individualised to be helpful, informing end-users about the RP data protection practices and the degree of compliance with their personal preferences.
- **Reduced complexity:** The provided recommendation must be meaningful and straightforward to avoid fatigue and obtain the desired results.
- **Considering the audience:** Showing all details regarding the recommendation at once is rarely practical. However, expert users will probably require an explanation of the provided recommendation. The recommendation must be given in layers, from the straightforward recommendation to different levels of detail coming from the Data Collection module and the Feature Collection module.
- **Offering meaningful choices:** The provided recommendation must be actionable, enabling users to make informed privacy decisions.

The first three characteristics must be ensured by how the recommendation is built and delivered, considering the profiles stored in the Feature Collection module and providing multilayered visual methods combining text, images and icons. Multilayered recommendations constitute a set of complementary details tailored to different audiences, carefully designed in terms of presentation and layout to present and extend information gradually to impact users' attention and comprehension of the recommendation. Colour is also essential; the PAdv relies on the traffic light colour code (green, yellow and red) and uses grey for unknown categories or attributes.

The fourth characteristic of the list, regarding offering meaningful choices, is ensured by the advisor's integration in the authentication and authorisation flows as explained in Section 3.3: the end-user can make critical decisions based on the received recommendation. Mainly, she can continue, limit or abort the interaction with the RP, or change her account configuration in the IdP.

4.3.1. First Layer

As introduced before, the first layer recommendation is delivered with a traffic light colour code. Green means OK (Good). Therefore, the end-user can continue the interaction with the RP with privacy guarantees considering her profile and the available information about that RP. Yellow means partially OK; therefore, a "Caution" recommendation is made. The end-user should limit interactions with the RP (only if there are no alternatives, only if strictly necessary) or change the account's configuration at the IdP to improve the control

over the PII sharing. Finally, red means NO OK (Bad) and a recommendation to abort the interaction with the RP. The grey colour is for “Unknown” results, meaning that it is impossible to provide a personalised or well-grounded recommendation due to lack of information.

As has been already mentioned, content-based filtering is used to obtain this recommendation, specifically based on decision trees. Different decision trees generate recommendations for different RP categories as the recommendation changes depending on whether the accessed RP is in the News, Finance or Health category, for example.

Figure 3 shows an example of the decision tree used for the E-commerce and Shopping category. The recommendation can be Unknown if there is not enough available data about the RP. Or Good, Caution or Bad depending on the Total Risk Score, used as the Rating in the decision tree (from 1 to 256). When changing from one category of RP to another, more or less the privacy-critical category, the only things that change in the decision trees are the R ranks for each recommendation.

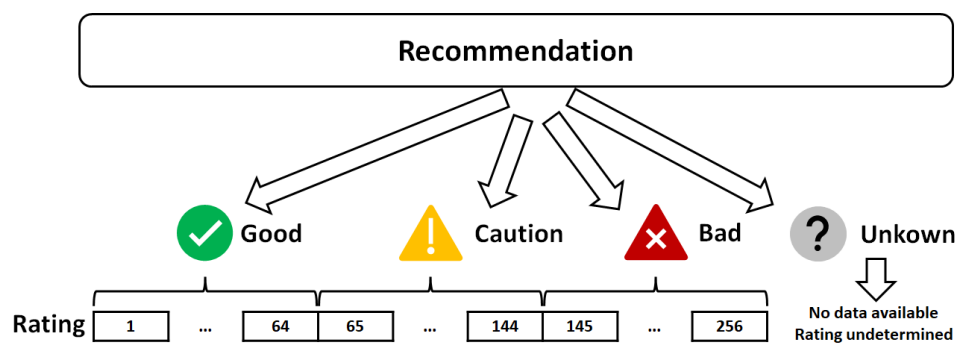


Figure 3. Example of decision tree for the E-commerce and Shopping category of RP.

Data collected with the Data Collection module are used to tag a specific recommendation request with a Transaction risk score. In this case, the risk is associated not with the end-user as in the case of Inherent risk, but with the use of a specific RP at a specific moment, hence the term “Transaction” risk.

Again, heat maps are used, in this case considering features of the assessed RP. Attributes regarding certifications and compliance are represented in one dimension (Table 3, 1: excellent, 2: good, 3: fair and 4: poor). They are combined into a single dimension because they eventually determine the same thing, conformity with a law, regulation, technical specification, certification associated with good practice, etc. Attributes regarding design patterns are represented in the other one (Table 4, 1: excellent, 2: good, 3: fair and 4: poor). The results obtained from the Reputation component, if available, are used to adjust the obtained score. For example, an RP with a good score in certifications and compliance that has not been found to incorporate dark design patterns, $2(\text{good}) \times 1(\text{excellent}) = 2$, may end up with a higher Transaction risk if the site’s reputation is terrible because it has suffered several data breaches in recent months or because it has been shown to profile its users aggressively.

Table 3. Certifications and compliance.

Score	Explanation
1—Excellent	At least one privacy certification and at least 90% compliant with checked best practices and regulation
2—Good	At least one privacy certification and at least 60% compliant with checked best practices and regulation
3—Fair	No privacy certification and at least 40% compliant with checked best practices and regulation
4—Poor	No privacy certification and below a 40% of compliance

Table 4. Design patterns.

Score	Explanation
1—Excellent	No identified dark patterns
2—Good	Using one dark pattern
3—Fair	Using between two and four dark patterns
4—Poor	Using more than four dark patterns

In particular, reputation is used to adjust the certification and compliance score. The performed static evaluation on the application of good practices must be calibrated based on what other users or independent agents have observed about how these practices behave in real use cases. Table 5 shows the proposed calibration. The score can in no case go below 1 or above 4: if a provider is already at the lowest or highest possible score, its bad or good reputation does not change this score any further, it does not add anything new to the performed assessment.

Table 5. Calibration of the certification and compliance score based on the reputation of the RP involved in the transaction.

Calibration	Explanation
0	The initial score is not modified because there is no available information about Reputation or because it is between a 40 and a 60 (intermediate or average values)
+1	The initial risk score goes up one level because Reputation is below a 40 (the RP has not good reputation)
−1	The initial risk score goes down one level because Reputation is above a 60 (the RP has good reputation)

The Total Risk Score used as the Rating in the recommender decision trees is computed as $R = \text{inherent risk} \times \text{transaction risk}$, ranging from 1 (1×1) to 256 (16×16).

4.3.2. Second Layer

The second layer shows the Inherent and Transaction risk scores for the user and the RP, respectively, as well as the used decision tree and the heat maps (Figure 4) for the user to understand how the recommendation has been generated. Suppose the recommendation is “Caution”, this second layer also lists the recommended actions for the end-user: how to restrict the interaction with the RP or how to change the account configuration at the IdP. These changes may modify the end-user profile (her exposure, for example, sharing less PII with the IdP) and, therefore, the recommendation in future interactions.

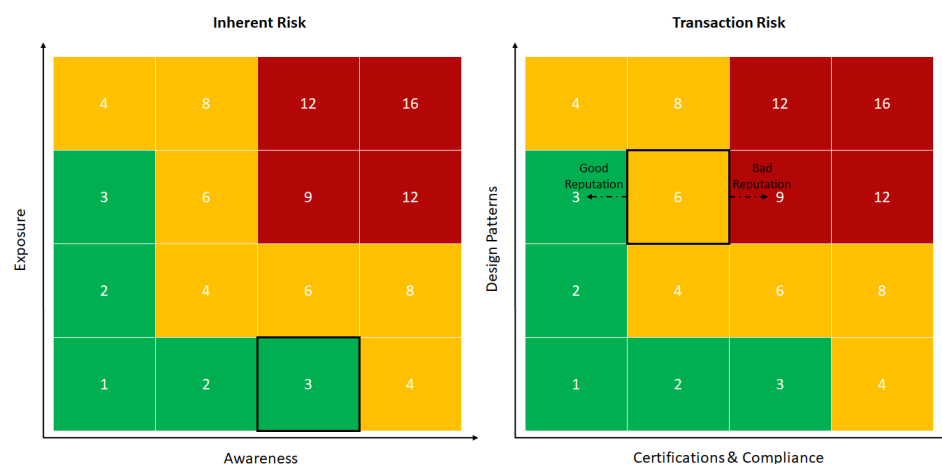


Figure 4. Example of heat maps for the Inherent and Transaction risks.

4.3.3. Third Layer

Finally, the third layer shows a drop-down menu with all the data collected by the Data collection component and the source or origin of each of them. Figure 5 shows an example of the information displayed in the three proposed recommendation layers.

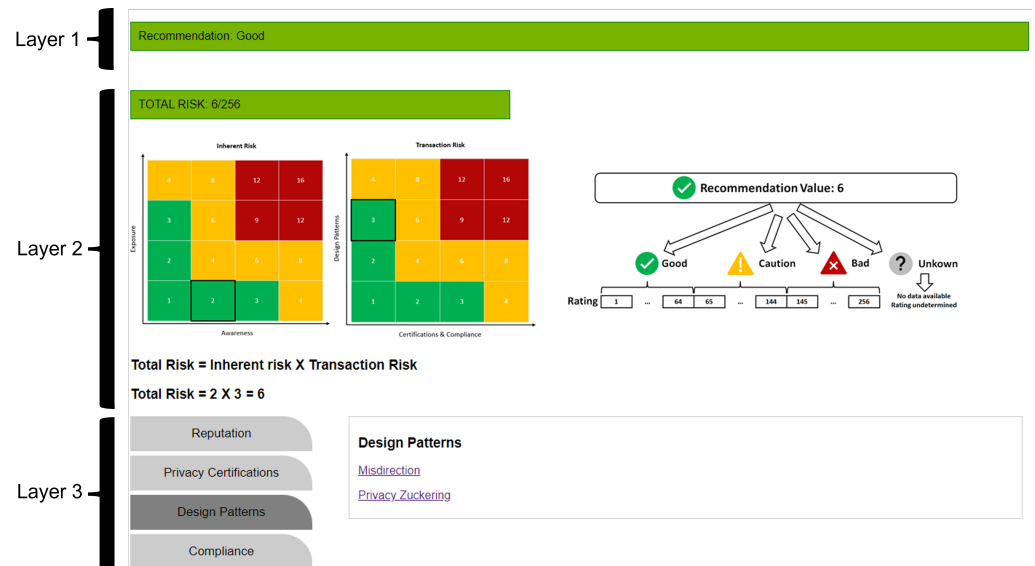


Figure 5. Summary of the information provided by the different PAdv recommendation layers.

5. Proof of Concept: Validating and Evaluating a Prototype of the Privacy Advisor

A new plugin has been developed for Google Chrome which opens a new tab when users click on an IdP sign-in link. This tab shows the privacy recommendation for the service the user wants to use. This implementation alternative has been selected because it is flexible and enables validation from different devices (laptops, smartphones, and tablets) interacting with different RPs.

The PAdv implementation is centralised in this proof of concept, using Python 3, running on a Windows server with an Intel Core i5 processor and 8 GB of RAM. The Recommendation Module receives an HTTPS request from the end-user, and, once the other modules have been consulted, it composes an HTML web page with all the information about the requested service to be presented to the user following the proposed multilayered approach already introduced.

The authentication flow executes independently, and it has been certified that it is not altered by this new entity, the Privacy Advisor.

A cache that stores the users’ requests helps the Privacy Advisor to reduce its response time. The response of every component is stored, associated with the requested RP. When a user asks for a recommendation, the Privacy Advisor searches if this RP has been previously assessed. The PAdv provides the information stored in the cache if it has been. This cache has 10 min expiration time to avoid massive storage and unnecessary resource consumption.

Each performed test asks the Privacy Advisor to generate a completely new recommendation for an RP within a subset of the most popular per SimilarWeb category. The results for each performance figure have been obtained by executing fifty experiments (twice per RP), and computing arithmetic means obtaining the following:

- Latency: The average time of a whole new recommendation (flow shown in Figure 1), without using the cache, is 16.1 s (with 510 ms of standard deviation). Again, on average, it takes 700 milliseconds to conduct communication tasks between modules and the rest of the time is consumed by the different modules obtaining their results from scratch (some of them require time-consuming tasks such as HTML parsing and analysis). Since end-users often use the same services, the latency could be significantly improved using the cache and proactively analysing some providers in advance

(for example, similar to the set of providers currently used by the end-user or other similar users). The average time to obtain a recommendation is 5.4 ms (with 0.15 ms of standard deviation) with these improvements.

- Resource consumption usage: During the Privacy Advisor’s execution, the use of resources on average is 30 MB RAM and 11% of CPU. The highest usage of RAM is 38 MB and 13% for the CPU.

A survey with 151 participants was carried out to prove the Privacy Advisor’s potential acceptance, considering its benefits, effectiveness and usability.

Regarding demographics, our sample population includes 65% of male participants and 33% of female participants (left part of Figure 6). All of them are from Spain. A large group works in the information technology sector, so it is expected that their awareness of privacy is slightly greater than in other sectors. The right part of Figure 6 shows the age of the participants. Additionally, Figure 7 shows the participants’ working sector grouped using The Global Industry Classification Standard [53]. The Others category includes other employment statuses such as students or unemployed participants. Only these data about the participants have been collected, gender, age and working sector. In this way, avoiding the gathering of personally identifiable information, the risks of the validation experiments are minimised, as they are always aimed at evaluating the proposed solution and its usefulness.

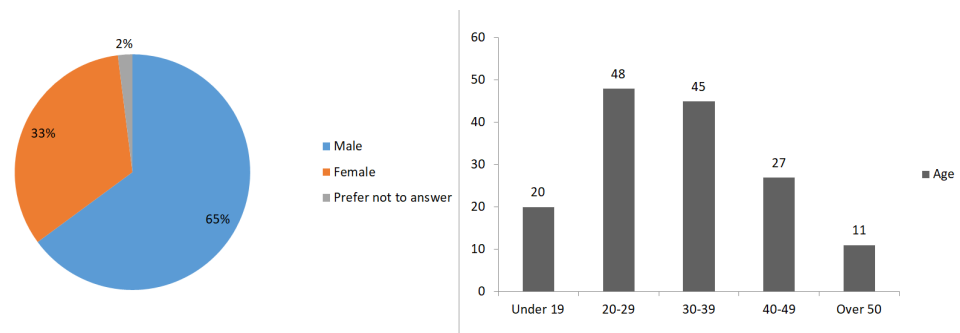


Figure 6. Participants gender (left) and age (right).




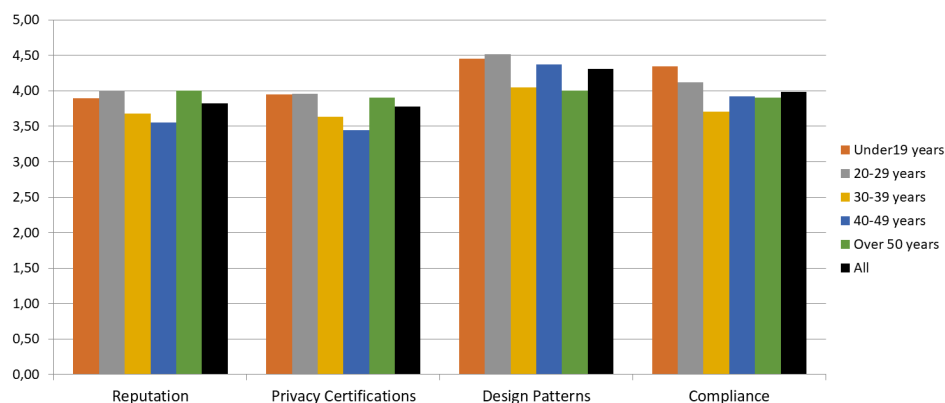
Figure 7. Participants working sector.

A questionnaire was prepared (Table 6) that asks about the quality of the recommendation provided for the 25 different RPs selected for the PAdv prototype evaluation, the quality of the decisions that have been made and the effort that these decisions have entailed for the user. Participants were asked to answer each question using stars and a Likert scale: Strongly disagree/Disagree/Neither agree nor disagree/Agree/Strongly agree.

Table 6. Questionnaire and average scores.

#	Question	Measurement	Average Score
1	The recommendations match your needs and are personalised	Recommendation quality	4.12
2	The recommendations are actionable and allow you to make decisions about your privacy protection that you would not have made without the PAdv	Recommendation quality	4.28
2	You are satisfied with your final choices	Decision quality	4.05
3	The decision processes have been easy	Decision effort	3.84
4	The time devoted to make decisions is affordable	Decision effort	3.23

The questions shown in Table 6 refer to the PAdv as a whole and to the recommendation it provides. A specific question to assess the utility of the different kinds of information provided in the PAdv's third layer was added to the questionnaire, allowing participants to rate the utility of the PAdv components: Design patterns component get 4.31 out of 5, Privacy certifications component get 3.77 out of 5, Compliance component get 3.99 out of 5 and Reputation component get 3.82 out of 5. Figure 8 shows this scoring in detail, revealing again that all components are perceived as practical and valuable by the different age groups.

**Figure 8.** Data collection components rating.

Discussion

End-users' privacy decisions may be involuntary, poorly founded or avoidable due to their knowledge or awareness limitations, time constraints, providers' opaque business models and incentives and ecosystems complexity. In this sense, the proposed model has proven to be a valuable tool to empower users in their relationship with different providers (of resources, applications, services, and identities). In an identity federation that includes a PAdv, users will be able to make better decisions about protecting their personal data and do so efficiently, compared to a federation in which this actor does not exist.

The proposed model may have some limitations from a design point of view. The first relates to the need to add a new agent or actor to a context that is itself already multi-agent or multi-actor identity federations. Whether this is possible in practice will depend on the availability of organisations that are in a position to operate reliable PAdv.

The second relates to the user-centric approach to privacy. Although the proposed recommender system enables end-users to make informed decisions about their privacy by providing a system that supports these decisions, users still need to have some knowledge or expertise to ensure proper privacy self-management. As previous work has discussed, user-centric approaches usually forget that the cause of the appropriation of users' personal data is not the failure of end-users when deciding about their data protection but business models based on surveillance [54,55].

Finally, there may be a technical limitation caused by the asymmetry of information sharing on the Internet and other digital ecosystems: the PAdv is based on a Data Collection module, but it can be challenging to obtain the necessary data on how an identity or resource

provider operates to make a recommendation of sufficient quality. Providers are regularly opaque about their data capture, processing, and protection practices.

Future work described in the following section will try to overcome some of these limitations.

6. Conclusions and Future Work

This paper proposes a new model based on a recommendation system, the Privacy Advisor, which is a decision support system to assist end-users when selecting service providers or choosing their privacy settings within identity federations. It can provide personalised privacy recommendations based on real-time data collection about the assessed providers. Implementing a Privacy Advisor prototype has allowed us to validate the proposed model and evaluate the effectiveness of our approach. A field study performed with 151 participants demonstrated the Privacy Advisor's utility in assisting users in their decisions. The results of this study are encouraging, showing substantial engagement with provided recommendations with good decision quality and decision support outcomes.

We are investigating how to fingerprint providers to collect helpful information on how they protect their users' privacy, even when their methods are opaque. On the one hand, we are investigating different techniques to collect information on the attributes already covered in this research (design patterns, privacy certifications, and compliance). On the other hand, by incorporating new attributes such as the designation of a data protection officer, the country in which the data are stored or the analysis of privacy policies. For the latter, we are exploring the use of natural language processing techniques.

We are also working on building provider reputation systems, both in the protocols that support them and in the definition of reputation metrics or the weights assigned to these metrics depending on the evaluator.

Similarly, we are making progress on hybrid recommendation systems that consider the user for whom the recommendation is made and other similar users. In other words, we are trying to merge Collaborative and Content-based filtering to combine the strengths of both approaches.

Finally, we would like to work on ways to encourage the collaboration of the different actors in identity federations to offer and use PAdv and reputation evaluators to be more transparent and make data collection more accessible.

Author Contributions: C.V.: Methodology, Software, Validation, Formal analysis, Investigation, Data Curation, and Writing—Original Draft. M.B.: Conceptualisation, Validation, Investigation, Writing—Review and Editing, and Supervision. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Ethical review and approval were waived for this study, due to the nature of the performed experiments and gathered data (not identifiable—gender, working sector and age- and involving a minimal risk).

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: No new data were created or analyzed in this study, only those related to the validation and evaluation of the implemented prototype. Data sharing is not applicable to this article.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Ahn, G.J.; Lam, J. Managing privacy preferences for federated identity management. In Proceedings of the 2005 Workshop on Digital Identity Management, Fairfax, VA, USA, 11 November 2005; pp. 28–36.
2. Barth, S.; de Jong, M.D.; Junger, M.; Hartel, P.H.; Roppelt, J.C. Putting the privacy paradox to the test: Online privacy and security behaviors among users with technical knowledge, privacy awareness, and financial resources. *Telemat. Inform.* **2019**, *41*, 55–69. [[CrossRef](#)]

3. RFC 6749: The OAuth 2.0 Authorization Framework. 2012. Available online: <https://tools.ietf.org/html/rfc6749> (accessed on 1 January 2022).
4. OpenID Connect Core 1.0 Incorporating Errata Set 1. 2014. Available online: https://openid.net/specs/openid-connect-core-1_0.html (accessed on 1 January 2022).
5. Mobile Connect. Available online: <https://mobileconnect.io> (accessed on 1 January 2022).
6. Murmann, P.; Karegar, F. From design requirements to effective privacy notifications: Empowering users of online services to make informed decisions. *Int. J. Hum. Comput. Interact.* **2021**, *37*, 1823–1848. [CrossRef]
7. Knijnenburg, B.P.; Willemsen, M.C.; Hirtbach, S. Receiving recommendations and providing feedback: The user-experience of a recommender system. In Proceedings of the International Conference on Electronic Commerce and Web Technologies, Munich, Germany, 1–4 September 2010; Springer: Berlin/Heidelberg, Germany, 2010; pp. 207–216.
8. Zhang, J.; Curley, S.P. Exploring Explanation Effects on Consumers' Trust in Online Recommender Agents. *Int. J. Hum. Comput. Interact.* **2018**, *34*, 421–432. [CrossRef]
9. Xiao, B.; Benbasat, I. An empirical examination of the influence of biased personalized product recommendations on consumers' decision making outcomes. *Decis. Support Syst.* **2018**, *110*, 46–57. [CrossRef]
10. Sonboli, N.; Smith, J.J.; Cabral Berenfus, F.; Burke, R.; Fiesler, C. Fairness and transparency in recommendation: The users' perspective. In Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization, Utrecht, The Netherlands, 21–25 June 2021; pp. 274–279.
11. Isaakidis, M.; Halpin, H.; Danezis, G. UnlimitID: Privacy-preserving federated identity management using algebraic MACs. In Proceedings of the ACM on Workshop on Privacy in the Electronic Society, Vienna, Austria, 24–28 October 2016; pp. 139–142.
12. Navas, J.; Beltrán, M. Understanding and mitigating OpenID Connect threats. *Comput. Secur.* **2019**, *84*, 1–16. [CrossRef]
13. del Álamo, J.M.; Monjas, M.; Yelmo, J.C.; González, B.S.M.; Trapero, R.; Fernández, A.M. Self-service Privacy: User-Centric Privacy for Network-Centric Identity. In Proceedings of the Trust Management IV—4th IFIP WG 11.11 International Conference, Morioka, Japan, 16–18 June 2010; IFIP Advances in Information and Communication Technology; Springer: Berlin/Heidelberg, Germany, 2010; Volume 321, pp. 17–31.
14. del Alamo, J.M.; Fernandez, A.M.; Trapero, R.; Yelmo, J.C.; Monjas, M.A. A Privacy—Considerate Framework for Identity Management in Mobile Services. *Mob. Networks Appl.* **2011**, *16*, 446–459. [CrossRef]
15. Sánchez, R.; Almenares, F.; Arias, P.; Díaz-Sánchez, D.; Marín, A. Enhancing Privacy and Dynamic Federation in IdM for Consumer Cloud Computing. *IEEE Trans. Consum. Electron.* **2012**, *58*, 95–103. [CrossRef]
16. Zwattendorfer, B.; Slamánig, D.; Stranacher, K.; Hörandner, F. A Federated Cloud Identity Broker-Model for Enhanced Privacy via Proxy Re-Encryption. In Proceedings of the 15th IFIP TC 6 TC 11 International Conference on Communications and Multimedia Security, Aveiro, Portugal, 25–26 September 2014; Lecture Notes in Computer Science; Springer: Berlin/Heidelberg, Germany, 2014; Volume 8735, pp. 92–103.
17. Asghar, M.R.; Backes, M.; Simeonovski, M. PRIMA: Privacy-preserving identity and access management at internet-scale. In Proceedings of the IEEE International Conference on Communications (ICC), Kansas City, MO, USA, 20–24 May 2018; IEEE: Piscataway, NJ, USA, 2018; pp. 1–6.
18. Henriksen-Bulmer, J. Incorporating Contextual Integrity into Privacy Decision Making: A Risk Based Approach. Ph.D. Thesis, Bournemouth University, Poole, UK, 2019.
19. Murmann, P.; Beckerle, M.; Fischer-Hübner, S.; Reinhardt, D. Reconciling the what, when and how of privacy notifications in fitness tracking scenarios. *Pervasive Mob. Comput.* **2021**, *77*, 101480. [CrossRef]
20. Alemany, J.; del Val, E.; Alberola, J.; García-Fornes, A. Enhancing the privacy risk awareness of teenagers in online social networks through soft-paternalism mechanisms. *Int. J. Hum. Comput. Stud.* **2019**, *129*, 27–40. [CrossRef]
21. Ghazinour, K.; Matwin, S.; Sokolova, M. *Monitoring and Recommending Privacy Settings in Social Networks*; EDBT/ICDT: Genoa, Italy, 2013.
22. Ghazinour, K.; Matwin, S.; Sokolova, M. YourPrivacyProtector: A Recommender System for Privacy Settings in Social Networks. *Int. J. Secur. Priv. Trust. Manag.* **2013**, *2*, 11–25. [CrossRef]
23. Zhang, Y.; Humbert, M.; Rahman, T.; Li, C.T.; Pang, J.; Backes, M. Tagvisor: A Privacy Advisor for Sharing Hashtags. In Proceedings of the WWW 2018: The 2018 Web Conference, Lyon, France, 23–27 April 2018.
24. Orekondy, T.; Schiele, B.; Fritz, M. Towards a Visual Privacy Advisor: Understanding and Predicting Privacy Risks in Images. In Proceedings of the International Conference on Computer Vision, Venice, Italy, 22–29 October 2017.
25. Chairani, M.; Chevalley, M.; Lazraq, A.; Bhagavatula, S. By the user, for the user: A user-centric approach to quantifying the privacy of websites. *arXiv* **2019**, arXiv:1911.05798.
26. Bernsmed, K.; Tøndel, I.A.; Nyre, Å.A. Design and Implementation of a CBR-based Privacy Agent. In Proceedings of the Seventh International Conference on Availability, Reliability and Security, Prague, Czech Republic, 20–24 August 2012; IEEE Computer Society: Washington, DC, USA, 2012; pp. 317–326.
27. Chang, C.; Li, H.; Zhang, Y.; Du, S.; Cao, H.; Zhu, H. Automated and personalized privacy policy extraction under GDPR consideration. In Proceedings of the International Conference on Wireless Algorithms, Systems, and Applications, Honolulu, HI, USA, 24–26 June 2019; Springer: Berlin/Heidelberg, Germany, 2019; pp. 43–54.
28. Liu, R.; Cao, J.; Zhang, K.; Gao, W.; Liang, J.; Yang, L. When privacy meets usability: Unobtrusive privacy permission recommendation system for mobile apps based on crowdsourcing. *IEEE Trans. Serv. Comput.* **2016**, *11*, 864–878. [CrossRef]

29. Liu, B.; Andersen, M.S.; Schaub, F.; Almuhiemedi, H.; Zhang, S.A.; Sadeh, N.; Agarwal, Y.; Acquisti, A. Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions. In Proceedings of the Twelfth Symposium on Usable Privacy and Security (SOUPS, Santa Clara, CA, USA, 12–14 July 2017; USENIX Association: Berkeley, CA, USA, 2016; pp. 27–41.
30. Andow, B.; Mahmud, S.Y.; Whitaker, J.; Enck, W.; Reaves, B.; Singh, K.; Egelman, S. Actions Speak Louder than Words: Entity-Sensitive Privacy Policy and Data Flow Analysis with PoliCheck. In Proceedings of the 29th USENIX Security Symposium, Boston, MA, USA, 12–14 August 2020; pp. 985–1002.
31. Sanchez, O.R.; Torre, I.; He, Y.; Knijnenburg, B.P. A recommendation approach for user privacy preferences in the fitness domain. *User Model. User Adapt. Interact.* **2019**, *30*, 513–565. [[CrossRef](#)]
32. Keshavarz, M.; Anwar, M. Towards Improving Privacy Control for Smart Homes: A Privacy Decision Framework. In Proceedings of the 16th Annual Conference on Privacy, Security and Trust (PST), Belfast, Ireland, 28–30 August 2018.
33. Resnick, P.; Varian, H.R. Recommender Systems. *Commun. ACM* **1997**, *40*, 56–58. [[CrossRef](#)]
34. Pawlicka, A.; Pawlicki, M.; Kozik, R.; Choraś, R.S. A Systematic Review of Recommender Systems and Their Applications in Cybersecurity. *Sensors* **2021**, *21*, 5248. [[CrossRef](#)] [[PubMed](#)]
35. Ko, H.; Lee, S.; Park, Y.; Choi, A. A Survey of Recommendation Systems: Recommendation Models, Techniques, and Application Fields. *Electronics* **2022**, *11*, 141. [[CrossRef](#)]
36. Srfi, M.; Oussous, A.; Ait Lahcen, A.; Moulina, S. Recommender systems based on collaborative filtering using review texts—A survey. *Information* **2020**, *11*, 317. [[CrossRef](#)]
37. Al Hassanieh, L.; Abou Jaoudeh, C.; Abdo, J.B.; Demerjian, J. Similarity measures for collaborative filtering recommender systems. In Proceedings of the 2018 IEEE Middle East and North Africa Communications Conference (MENACOMM), Jounieh, Lebanon, 18–20 April 2018; pp. 1–5.
38. Javed, U.; Shaukat Dar, K.; Hameed, I.; Iqbal, F.; Mahboob Alam, T.; Luo, S. A Review of Content-Based and Context-Based Recommendation Systems. *Int. J. Emerg. Technol. Learn. (ijET)* **2021**, *16*, 274–306. [[CrossRef](#)]
39. Karimi, M.; Jannach, D.; Jugovac, M. News recommender systems—Survey and roads ahead. *Inf. Process. Manag.* **2018**, *54*, 1203–1227. [[CrossRef](#)]
40. Singh, P.K.; Pramanik, P.K.D.; Dey, A.K.; Choudhury, P. Recommender systems: An overview, research trends, and future directions. *Int. J. Bus. Syst. Res.* **2021**, *15*, 14–52. [[CrossRef](#)]
41. Top Websites Ranking—SimilarWeb. Available online: <https://www.similarweb.com/top-websites/> (accessed on 1 January 2022).
42. NIST Special Publication (SP) 800-30, Revision 1, Guide for Conducting Risk Assessments. 2020. Available online: <https://www.nist.gov/privacy-framework/nist-sp-800-30> (accessed on 1 January 2022).
43. Yang, Y.; Du, X.; Yang, Z. PRADroid: Privacy Risk Assessment for Android Applications. In Proceedings of the 5th IEEE International Conference on Cryptography, Security and Privacy, CSP, Zhuhai, China, 8–10 January 2021; IEEE: Piscataway, NJ, USA, 2021; pp. 90–95.
44. Dark Patterns. Available online: <https://www.darkpatterns.org/> (accessed on 1 January 2022).
45. European Privacy Seal (EurPriSe). Available online: <https://www.european-privacy-seal.eu> (accessed on 1 January 2022).
46. TrustArc. Available online: <https://www.trustarc.com> (accessed on 1 January 2022).
47. ePrivacy. Available online: <https://www.eprivacy.eu/> (accessed on 1 January 2022).
48. Jaithunbi, A.; Sabena, S.; SaiRamesh, L. Trust evaluation of public cloud service providers using genetic algorithm with intelligent rules. *Wirel. Pers. Commun.* **2021**, *121*, 3281–3295. [[CrossRef](#)]
49. Kokoris-Kogias, E.; Voutyras, O.; Varvarigou, T. TRM-SIoT: A scalable hybrid trust & reputation model for the social Internet of Things. In Proceedings of the 2016 IEEE 21st International Conference on Emerging Technologies and Factory Automation (ETFA), Berlin, Germany, 6–9 September 2016; pp. 1–9.
50. Barakat, L.; Taylor, P.; Griffiths, N.; Miles, S. A Reputation-based Framework for Honest Provenance Reporting. *ACM Trans. Internet Technol.* **2021**. Available online: https://kclpure.kcl.ac.uk/portal/files/166020608/BarakatEtAl_1_.pdf (accessed on 1 January 2022).
51. Govindaraj, R.; Govindaraj, P.; Chowdhury, S.; Kim, D.; Tran, D.T.; Le, A.N. A Review on Various Applications of Reputation Based Trust Management. *Int. J. Interact. Mob. Technol.* **2021**, *15*, 87–102.
52. Zhou, Z.; Wang, M.; Yang, C.N.; Fu, Z.; Sun, X.; Wu, Q.J. Blockchain-based decentralized reputation system in E-commerce environment. *Future Gener. Comput. Syst.* **2021**, *124*, 155–167. [[CrossRef](#)]
53. The Global Industry Classification Standard (GICS). Available online: <https://www.msci.com/gics> (accessed on 1 January 2022).
54. Janssen, H.; Cobbe, J.; Singh, J. Personal information management systems: A user-centric privacy utopia? *Internet Policy Rev.* **2020**, *9*, 1–25. [[CrossRef](#)]
55. Kröger, J.L.; Lutz, O.H.M.; Ullrich, S. The Myth of Individual Control: Mapping the Limitations of Privacy Self-Management. Available online: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3881776 (accessed on 1 January 2022).