ARTICLE IN PRESS

Future Generation Computer Systems xxx (xxxx) xxx

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Semantic-based privacy settings negotiation and management

Odnan Ref Sanchez [a], Ilaria Torre [a,*], Bart P. Knijnenburg [b]

[a] Department of Informatics, Bioengineering, Robotics and Systems Engineering, University of Genoa, Genoa, Italy
[b] School of Computing, Clemson University, Clemson, SC, USA

ARTICLE INFO

Article history:
Received 19 July 2018
Received in revised form 8 August 2019
Accepted 27 October 2019
Available online xxxx

Keywords:
Privacy Preference Manager
Privacy ontology
Data sharing and permission management
IoT

ABSTRACT

By 2020, an individual is expected to own an average of 6.58 devices that share and integrate a wealth of personal user data. The management of privacy preferences across these devices is a complex task for which users are ill-equipped, which increases privacy risks. In this paper we propose an approach that exploits Semantic Web (SW) technology to manage the user's IoT privacy preferences and negotiate the permissions for data sharing with third parties. SW technology comprises a web of data that can be processed by machines through a formal, universally shared representation. In our approach, SW enables a lightweight and interoperable communication between a Personal Data Manager (PDM) and the Third Parties (TPs) that request access to the user's personal data. The PDM can handle multiple heterogeneous personal IoT devices and manages the negotiation process between the user and the TPs in a way that can relieve users from the burden of specifying their privacy requirement for each TP. The core of the approach is the definition of the Privacy Preference for IoT (PPIoT) Ontology which is based on the Privacy Preference Ontology, the W3C Semantic Sensor Network Ontology, the Fair Information Practices (FIP) principles, and state-of-the-art recommendation techniques for privacy protection in the IoT. This ontology aims to capture the complexity of privacy management in the IoT paradigm in light of the recent General Data Protection Regulation (GDPR) of the European Union. Along with presenting the ontology, in this paper we will provide an example on how to use the PPIoT ontology for the management of privacy preferences in the fitness IoT domain and we will show how the PDM handles the process of negotiation between the user and the TPs. The approach is based on an interactive PPIoT-based Privacy Preference Model (PPM) that meets the requirements of the GDPR to have transparent and simple TP privacy policies. Finally, we will report the results of an evaluation on a mockup fitness app that implements this PPM. The main contributions of this paper are: (i) to propose an ontology for privacy preference in the IoT context, which covers a knowledge gap in existing literature and can be used for IoT privacy management, (ii) to propose an interactive PPIoT-based Privacy Preference Model, which is in accordance with the GDPR objectives.

1. Introduction

The Internet of Things (IoT) is a network that gives physical devices – ranging from small sensors, personal devices, to larger devices such as smart TVs and smart connected cars – the ability to transmit and receive data [1]. The IoT will likely reach a combined total of 18 billion connections by 2022 [2]. Today, an individual owns an average of 3.64 connected devices, which is expected to grow to 6.58 devices by 2020 [3]. The rise of this technological paradigm has numerous advantages. However, these devices also pose threats to the user's privacy—especially those devices that track users' round-the-clock activities (e.g., fitness trackers). Indeed, the concern for online privacy management is exacerbated by the introduction of such IoT devices, which produce and share an enormous amount of personal data in addition to that on social systems [4,5]—meeting the criteria of Big Data in terms of volume, velocity, variety, veracity, and value [6]. The management of this data is a complex task, as it requires heterogeneous devices and applications to be managed accordingly [7].

In this paper, we propose an approach that exploits Semantic Web Technology (SWT) to manage the interactive setting of users' privacy preferences in the IoT. The core of the approach is a Privacy Preference Model (PPM) with its related ontology (PPIoT), and a Personal Data Manager (PDM) that negotiates and manages the user's privacy preferences.

The Semantic Web approach envisions a web of data that can be processed by machines through a formal, universally shared representation. This aim fits the IoT vision of creating an interoperable environment for devices. SWT provides enhanced services

* Corresponding author.
E-mail addresses: odnan.ref.sanchez@edu.unige.it (O.R. Sanchez), ilaria.torre@unige.it (I. Torre), bartk@clemson.edu (B.P. Knijnenburg).

https://doi.org/10.1016/j.future.2019.10.024
0167-739X/© 2019 Elsevier B.V. All rights reserved.

at the application layer, and is therefore considered a promising approach to manage interoperability among IoT silos [8,9]; the use of shared vocabularies and ontologies allows devices and services to communicate with each other, independently of their underlying implementation.

In our approach, we use SWT to formally describe both the user privacy preferences and the requirements of Third Parties (TPs) that request access to the user's personal data. Moreover, we use SWT to enable a lightweight and interoperable communication between the TPs and the PDM that is in charge of managing the user's privacy preferences.

For the representation of privacy preferences and for enabling the interaction between the PDM and the TPs, we created the Privacy Preference for IoT (PPIoT) Ontology (available online[1]). An ontology for representing users' privacy preferences has already been defined in [10], but that ontology models privacy preferences and access restrictions with a focus on social network applications. In contrast, our ontology has been created to capture privacy preferences specifically for the complex IoT paradigm and integrates/extends existing ontologies and state-of-the-art recommendation techniques to promote interoperability. The PPIoT ontology also follows the Fair Information Practices (FIP) principles [11], which are long-standing guidelines regarding the collection and use of users' information that aim to protect their privacy, and the General Data Protection Regulation (GDPR) [12], the European Regulation in force since May 25, 2018, to protect individuals against the processing and free movement of personal data.

In this paper we present the PPIoT ontology and provide an example on how to utilize it to give users more fine-grained control over their privacy preferences. We also provide a use case scenario showing the PDM capabilities for managing user privacy preferences, for negotiating between users and TPs, and for making privacy recommendations to the user. The main contribution of this paper is to address the current lack of solutions for the semantic management of user privacy in the field of IoT. Thanks to the Semantic Web layer, the PDM allows machine-driven negotiation among TPs on a universal set of parameters. It provides a uniform way of managing privacy preferences for heterogeneous TPs in a lightweight and scalable manner.

The GDPR also includes requirements of transparency and control regarding the communication of privacy information to individuals, aimed at countering complicated and lengthy "terms and conditions" which often include implicit consent for various data collection practices [13,14]. Based on the PPIoT ontology and the requirements above, we defined a Privacy Preference Model (PPM) that can be used by TPs to present their privacy policies to the user. The goal of PPM is to provide an interactive, simple, and straightforward way of presenting privacy policies to users, who can in turn give their affirmative consent for the data collection practices embedded in these policies in accordance to the GDPR [14].

We evaluate our PPIoT-based PPM with a mockup called "FitPro", which uses the PPM to present users its privacy policy and to allow them to accept or reject its various data collection practices. We recruited a total of 310 Fitbit fitness tracker users from the Amazon Mechanical Turk crowd-sourcing platform. They performed a simulated installation of the FitPro mockup and submitted a questionnaire asking them for feedback regarding the understandability of the app's privacy policy, the difficulty of setting their privacy permissions in the app, and the preferability of the PPM-based policy over traditional privacy policy presentations. The results of this evaluation confirm that users appreciate the interactive setting and controllability of privacy preferences. Furthermore, we compared participants' PPM-based settings against their existing real-world privacy settings, finding that they are significantly associated.

The remainder of this paper is structured as follows: Section 2 presents related work on privacy ontologies, privacy management, FIP, and GDPR. Section 3 explains the proposed PPIoT ontology and provides examples on how it is used to manage users' privacy preferences and TPs' privacy policy statements. Section 4.1 presents the PPM model, gives an overview of the PDM for interactive negotiation and recommendation, and elaborates on how the PDM manages the user's preferences, the negotiation between the user and TPs, and potential privacy recommendations to the user. A user evaluation of the PPM is presented in Section 5. Limitations of our approach and future work are discussed in Section 6, before concluding the paper in Section 7.

2. Background and related work

Previous work has already studied user privacy preference modeling, both with and without support of the SWT. P3P (the W3C Platform for Privacy Preferences) can be considered a reference model for the automatic processing of privacy preferences. Users can express their preferences, and their browser warns them if a site does not meet these preferences [15]. Many proposals about privacy enhancing technologies are based on P3P, which itself did not find a widespread application due to usability issues and a lack of enforcement. Moreover, the advent of social networks and the IoT brought new privacy requirements that are not implemented in P3P.

Using interviews and online surveys to model the privacy preferences of potential IoT users, Lee and Kobsa identified the contextual parameters that have the strongest influence on the user's privacy preferences [16]. These parameters include the type of monitoring, the type of information collected, the entity collecting the information, the frequency of monitoring, the location, and the reason for the collected data. Based on this data, Bahirat et al. created a privacy-setting interface that allows users to deny/allow IoT devices access to their personal information [7]. They also modeled users' decisions as a means to come up with default privacy profiles.

In this paper we propose an ontology that is specifically targeted to privacy management for IoT. Privacy preferences in the IoT context critically depend on the reason for data collection, the persistence of access, the location, the retention period and the method of usage [16–19]. These aspects constitute the requirements for privacy management in the IoT paradigm [20,21] and are taken into account in the proposed PPIoT Ontology. They also ascertain that PPIoT is compliant with the FIP principles, which require the requesting entity to clearly specify the reason, usage of data, frequency and method of data collection, and the retention period of the collected data.

Below, we first describe the base ontologies that are extended by our PPIoT ontology, and then we discuss other related ontologies for privacy modeling. After that, we present approaches for autonomous negotiation and discuss the FIP privacy principles and the GDPR regulation. This last section includes comparisons with other ontologies designed for, or addressing, the GDPR requirements.

2.1. Base ontologies

In line with best practices for ontology reuse, our PPIoT ontology integrates the current Privacy Preference Ontology (PPO) and the W3C Semantic Sensor Network (SOSA/SSN) Ontology. The PPO was created to aid users in managing their privacy settings in

[1] http://pdm-aids.dibris.unige.it/PPIoT

the realm of linked data [10]. Later, it was extended to facilitate social network applications [10]. The PPO allows users to have more fine-grained control over their personal data. One of the main features of this ontology is the possibility to set multiple privacy preferences for a user.

The original W3C Semantic Sensor Network (SSN) Ontology was aligned to the DOLCE-UltraLite3 Ontology and was based on the core concepts of the Stimulus–Sensor–Observation ontology pattern [22]. Due to the rapid expansion and diversity of data and its providers, it has been improved and is now based on the Sensor, Observation, Sample, and Actuator (SOSA) ontology pattern [23] to include broadened definitions (e.g., social sensing applications). As of October 2017, it became a W3C Recommendation.

2.2. Ontologies for privacy modeling

A survey paper by Perera et al. describes the history and state-of-the-art of ontology-based privacy modeling [24]. Below, we briefly describe the ontologies that are most related to our study.

The ontology described by Zhang and Todd shares similarities with PPO and is focused on defining privacy rules [25]. Each rule must contain a data class and a conditions class. The main features of the conditions class include the duration, purposes, and recipients of collection, how long the collected data will be retained, the user's privileges, and ways of handling disputes. This ontology was intended for applications of context-aware systems.

PROACT is an ontology that models privacy in relation to tasks and user activities [18]. The authors introduce the concept of an "activity sphere", which is a temporary abstract space defined to limit the incoming and outgoing information. PROACT is used to define privacy policies based on restrictions and rules for accessing and using each resource within an activity sphere and the information the resource collects and manages.

The authors of the PPO also created a light-weight ontology, named Privacy Preference Manager, which is a semantic representation of a tool that allows users to deny/allow access to their data based on the Web Access Control (WAC) vocabulary [10].

The privacy preference model proposed in Bodorik et al. is characterized by regulations and conditions that are specified by the user [19]. Conditions can concern the purpose of the data recipient, usage and retention, disputes, remedy, and access control. Bodorik et al. also specify the properties of a steady set of user preferences and their maintenance operations.

Privacy rules are also defined in the ontology by Hu and Yang [17]. Their ontology, which also follows the FIP guidelines, is intended to capture allowed/denied purposes of data collection, allowed/denied access for individual entities, retention period, obligations, policy, and action. This enables the creation of global rules to define preferences for higher-level conditions, such as giving a recipient access to data that came from medical applications, even if some of the needed parameters are not defined.

Setting rule priorities is another feature of privacy preference management. Rei is a policy language that aids users in expressing their privacy preference conditions with a priority hierarchy [26]. This level of expressiveness is a step forward in the enrichment of privacy specifications, as it helps in resolving conflicts. This is especially relevant in the context of IoT, where several conflicting conditions can occur. However, Rei still lacks the power of negotiation, since it can only set multiple conditions on the user side (i.e., it does not consider the TP side).

The usefulness of modeling trust is thoroughly described by Iqbal et al. [27]. Martimiano et al. model the trustworthiness of TPs using principles similar to Friend Of A Friend (FOAF) [28]. They extend this principle by defining fixed sets of classes with predefined assignments on the level of trust (e.g., close family, friend, work mates, unknown).

The ontologies mentioned in this subsection contain unique features that allow users to be more expressive regarding their privacy settings. However, these ontologies are limited in their application domain and focus exclusively on the user side, not on the TPs. To extend the ontological approach to privacy management to the field of IoT, it must have room for negotiation between the user and the TPs [24]. Our proposed PPIoT takes this into account.

A similar perspective is provided in [29], where the authors propose an approach to match users' privacy requirements against TPs' privacy statement. Similarly to our current work, they created a GDPR-based ontology, QoP, to support this match (it will be described in Section 2.4). Their approach supports the negotiation from the user side and, in this respect, it shares similarities with the part of our model for the interaction with ordinary TPs that do not have SWT capabilities and do not perform privacy negotiation.

2.3. Privacy negotiation

Trust-based negotiation has long been studied in literature given the nature of unknown recipients and benefactors of data (e.g., [30]). An early negotiation proposal is APPEL,[2] a language that extends P3P and enables users to express preferences as rules which can then be used by the user agent to make decisions and negotiation regarding privacy policies. For instance, Bennick et al. [31] utilize P3P/APPEL and propose basic negotiation mechanisms such as conditional accept and proposal reject that can be used during an isolated negotiation between users and TPs. Li et al. [32] also utilize the P3P/APPEL for user privacy modeling and provide two algorithms for negotiation. They propose Pareto optimal solutions and another algorithm that guarantees agreement after proposal exchanges between users and TPs. The combination of these algorithms provides seamless negotiation.

Ontology-based negotiation has also been proposed in the literature. Jang and Yoo [33] propose to quantify privacy sensitivity levels from unified personal information which they then use for the negotiation process. Jang et al. [34] also propose a negotiation system that mediates among the users, the service providers, and the law. Negotiation is done by taking into account the user's privacy preference and matching it to the TP privacy policy. Both studies use an ontology-based negotiation scheme.

Despite recent advances in privacy negotiation, there is no current standard that lets users attain agreement on privacy practices for IoT applications [35]. For this reason, proposing studies for IoT privacy negotiation is of interest for researchers.

A framework for privacy negotiation regarding Bluetooth Low Energy (BLE)-based IoT devices is implemented by Cha et al. [35]. As BLE devices do not have proper privacy policies towards user access, their *PrivacyBat* Framework of Privacy Preferences Expression for BLE-based applications allows user-side negotiation regarding nearby BLE devices. The framework negotiates with the user the specifications to achieve agreements on privacy practices. However, although it improves the privacy policy negotiation in IoT, users can essentially only accept or reject the policy.

Another study from Cha et al. [36] for privacy negotiation uses a Blockchain-Connected Gateway (BCG). The BCG acts as a mediator between users and IoT devices, allowing users to access device information and control. On the device side, all information towards the user will only be available if the user accepts

[2] W3C Working Draft: A P3P Preference Exchange Language (APPEL): https://www.w3.org/TR/P3P-preferences/

the privacy policy. BCGs are considered tamper-resistant, which can protect users when providing personal data to IoT devices. They can also store user privacy preferences on IoT devices in the blockchain network, which could resolve potential disputes between users and IoT service providers. While this approach resolves the storage of privacy preferences as another privacy issue, the negotiation still has only static accept or reject options.

A more dynamic negotiation framework can be found in Aydougan et al. [37]. In this study, aside from the static accept or reject options, users can negotiate with the TPs by excluding some parts of the requested data, or by asking for some (or a different type of) incentive. For this negotiation to work, the TP must have specific goals and a purpose for accessing the user data, and the user must have a motivation for sharing her personal data despite her privacy concerns. The approach is domain-dependent, as it needs to model the incentive and the user's type of information, which makes negotiation a rather complex ordeal.

Alanezi et al. [38] provide a framework for negotiation between the IoT user and the IoT deployment owner towards accessing an IoT service. The IoT owner is the responsible party for setting up and maintaining the IoT infrastructure that provides services to IoT users. The negotiation protocol uses XML to specify the privacy requirements of both parties. The negotiation algorithm sends a counter-proposal to the user if the accessing user's privacy policy does not match the owner's. However, while this solution allows the owner's side to negotiate, it is static on the user side.

A novel agent-based approach for negotiation is proposed in Baarslag et al. [39]. The negotiation is managed by an agent, based on the privacy preferences of actual users taken under different conditions. Although the agent effectively negotiates on the user's behalf, the researchers found that users still would like to engage in the negotiation process, as they usually do not trust automated processes regarding privacy. For this reason, the approach presented in the current paper always asks users themselves to make the final decision. Also, our framework provides both the user and enhanced TPs with the capability to negotiate (in the following we will explain that an "enhanced TP" is a TP with SWT capabilities).

2.4. Privacy principles, regulation and GDPR ontologies

Our PPIoT has been created following the principles of the long-standing Fair Information Practices (FIP) and the General Data Protection Regulation (GDPR). The FIP principles are conventional guidelines regarding the collection and use of users' information that aim to protect their privacy [40]. They include transparency, individual participation, purpose specification, data minimization, use limitation, data quality and integrity, security, accountability and auditing. The inclusion of the FIP principles in privacy frameworks is much needed, especially in the management of IoT user data collection [20,21,41], considering recent reports about the increase of privacy breaches of supposedly trusted TPs [41]. Hence, we have made sure that the PPIoT Ontology follows the FIP principles.

The GDPR is a legal framework of the EU enforced as of May 25, 2018 [14]. The principles relating to personal data protection are explicitly stated in Article 5 of the GDPR Regulation: lawfulness, fairness and transparency; purpose limitation; data minimization; accuracy; storage limitation; integrity and confidentiality. Broadly speaking, the GDPR requirements concern two main issues:
– management of personal data from the TP (data processing, sharing and storage),
– communication between TP and user about the management of personal data (transparency, controllability, accuracy).

The regulation applies even to TPs that have been established outside the EU, as long as they operate in an EU market or process the data of EU residents. It requires TPs to provide easily accessible and understandable privacy policies that use clear and plain language. Moreover, the user has to provide explicit consent to the privacy options expressed in these policies.

2.4.1. Concepts and terminology

Our PPIoT ontology is designed to include classes and properties that address the GDPR requirements for the management of personal data. Our focus is on the management of the user's privacy settings, not on TP obligations. Thus, the PPIoT ontology addresses only concepts related to users' decisions about privacy options in the TP statement. Below, we briefly describe the main concepts from the GDPR used in the PPIoT ontology and their mappings to the terms in the ontology (the PPIoT ontology will be described in detail in the next section):
– *Data subject* (Art. 4): corresponds to the *User* class;
– *Personal data* (Art. 4): corresponds to the *Dataset* class borrowed from the PPO ontology;
– *Controller*, *Processor*, *Third Party*, *Recipient* (Art. 4): all of these concepts are addressed with the *Entity* class, since their distinction is not relevant to the aims of a user-side privacy manager;
– *Consent* (Art. 4): in the GDPR, consent is defined as any freely given, specific, informed and unambiguous indication of the data subject's wishes signifying his/her agreement to the processing of personal data relating to him or her; as such, it corresponds to the final "Allow TP request" in our Interaction workflow (see Fig. 5), not to a concept in the ontology;
– *Processing* (Art. 4): is defined as any operation or set of operations which is performed on personal data, such as collection, recording, structuring, storage, dissemination or otherwise making available, etc.; since each type of processing has different properties, we modeled different types of processing separately, focusing, at the moment, on those that are highly relevant in a ubiquitous context: collection, storage and sharing. They are addressed by the properties *hasPersistency*, *hasMaxRetentionPeriod* and *allowsSharingWith*, respectively.
– *Purpose* (Art. 5): matches the *Reason* class and *hasReason* property, meaning that the consent must be bound to one or several specified purposes;
– *Conditions for consent* (Art. 7): are addressed by the *Condition* class from the PPO ontology; it contains the privacy conditions for the user's consent—which will be provided to a specific TP request;
– *Security of personal data* (Art. 32): is addressed within the *Method* class and the *hasMethod* property. Methods to secure personal data are particularly relevant in the IoT domain and are therefore usually seen as a precondition for the user to give her/his consent;

In addition, our interactive PPIoT-based Privacy Preference Model (PPM) conforms to GDPR requirements for the communication between the TP and the user, addressing transparency, controllability and accuracy issues.

2.4.2. GDPR-related ontologies

Since the adoption of the GDPR, several ontologies have been designed to support machine-processable reference to it and to address its requirements.

GDPRtEXT[3] [42] is an ontology designed to provide a way to refer to the concepts and terms expressed within the GDPR. It uses the ELI (European Legislation Identifier) OWL ontology to refer to different resources within the GDPR in terms of chapters,

[3] https://w3id.org/GDPRtEXT

Please cite this article as: O.R. Sanchez, I. Torre and B.P. Knijnenburg, Semantic-based privacy settings negotiation and management, Future Generation Computer Systems (2019), https://doi.org/10.1016/j.future.2019.10.024.

sections, articles, points and sub-points of the GDPR text. Moreover, it uses the W3C SKOS vocabulary to provide descriptions of GDPR concepts. This ontology does not aim to provide an interpretation of compliance obligations. Instead, it represents GDPR text as a set of RDF resources. Since the ontology allows one to refer to specific parts and concepts of the GDPR text, it is complementary to (rather than competing with) our PPIoT. After version 1 is released, we plan to provide alignments to this ontology, by referring privacy conditions to the exact articles of the GDPR text and by formally specifying the meaning of concepts (e.g., retention of personal data refers to *RetentionOfPersonalData* class which is defined through the property *involves*, and the *PersonalData* and *StoreData* classes).

Nonetheless, PPIoT does not need to refer to all of the GDPR concepts and obligations: it addresses only those obligations that are related to the expression of the user's consent and to negotiation with third party entities. This excludes, for instance, obligations to secure personal data and to maintain records of processing activities, which are imposed by law independently of the users' preferences. With regard to compliance checks of the entire set of GDPR obligations, tools have been developed for self-assessment (see, for instance, those developed by the Information Commissioner's Office in the UK [43] and by Microsoft [44]). In this respect, [45] proposes an extension of the W3C Open Digital Rights Language (ODRL) with the aim to represent both digital rights and legislative obligations, and applies the approach to GDPR through the development of a compliance assessment tool.

Besides GDPRtEXT, another recent legal ontology on GDPR is PrOnto [46]. Its goal is to provide legal knowledge modeling of GDPR concepts such as privacy agents, data types, types of processing operations, rights and obligations, with the final aim of supporting legal reasoning and checking compliance [47]. In the authors' own words, the ontology is still a draft, but useful alignments with our PPIoT could be made, should PrOnto be developed further. For example, PrOnto's *Data* sub-classes could extend our *Dataset* class, by specifying the type of data as being personal data, non-personal data, anonymized data, or pseudonymized data. Another example could be the alignment between the PPIoT *Reason* class and the PrOnto *Purpose* class and sub-classes that specify the reason for collecting and processing user data.

In addition to the ontologies mentioned above, which represent GDPR concepts independently from an application task, other ontologies have been designed with more focused goals.

Following the GDPR implementation, Elluri et al. [48] developed an ontology to represent some GDPR rules that concern Cloud data. The ontology is focused on the obligations of both the cloud data consumer and provider.

With respect to provenance modeling, GDPRprov [49] is an ontology aimed at modeling provenance for GDPR compliance. This is of course related to our approach, but it concerns the TP's task of recording the origin of data concerning the obtained consent and tracking its use and changes over time.

Finally, we already mentioned in Section 2.2 the Quality of Protection (QoP) ontology used in [29] to compare the users' privacy requirements against the TP Terms of Service (i.e., the TP policy statement). The paper presents only the main classes and subclasses that are used for the matching task. With respect to our ontology, it is not focused on user preferences, nor on the IoT domain.

In the following section, we present the PPIoT ontology, while the subsequent section describes the Privacy Preference Model for interactive privacy setting.

3. The privacy preference for IoT ontology

The guiding principle for the design of the PPIoT ontology is that it should be able to on the one hand represent the conditions of the user's privacy preferences regarding certain personal data (the Privacy Preference class), and on the other hand the conditions of the privacy policy statement of a TP entity (the Statement class). Thus, the goal is that both the user and the TP, through their respective applications, can set conditions regarding the access to the user's personal data (the Dataset class) that are produced by the IoT devices.

The PPIoT ontology is shown in Fig. 1. It has been designed to extend well-established existing ontologies with the aim of interoperability. In Fig. 1, existing ontologies are represented by the green, orange and blue nodes, while the black nodes are the extensions that we propose to cater for the privacy management needs in the context of the IoT.

3.1. Main imported ontology classes and properties

The core ontologies consist of PPO,[4] and SOSA/SSN.[5] Only the most relevant classes and properties that are imported from these ontologies are described in this section. Other ontologies that are used include FOAF[6] ACL,[7] WO,[8] VOID[9] and XSD[10] for defining data types. To avoid ambiguity, in this section a "statement" means an RDF statement while "[privacy] [policy] Statement" is the TP Statement that contains all the details about the data access request. The main imported Classes of the PPIoT are described below.

- ppo:PrivacyPreference: contains the user's privacy preferences defined in terms of conditions on personal user data (void:Dataset);
- ppo:Condition: contains conditions defined in terms of properties that denote restrictions to a specific Dataset instance;
- ppo:Operator: is a class having logical operators as subclasses (i.e., ppo:Or, ppo:And and ppo:Not) to allow more expressive conditions;
- sosa:Platform: any entity that hosts other entities, actuators, sensors, samplers, and even other platforms. Given the extensiveness of this class, we added a subclass for IoT devices;
- sosa:Sensor: a device, an agent (including humans), or software (simulation) involved in or implementing a procedure. Sensors respond to a stimulus (e.g., a change in the environment, or input data from the results of prior observations) and generate a result;
- wo:Weight: is a class defining values that specify the priority (rank) of a privacy preference;
- acl:Access: any kind of access mechanism to a resource;
- foaf:Agent: an agent (e.g., person, group, software or physical artifact);
- void:Dataset: the type of datasets that are collected, generated, maintained, or aggregated by an entity (e.g., user name, activity, weight, hearth rate, etc.).

The main imported Properties are:

[4] ppo: http://vocab.deri.ie/ppo#.

[5] ssn: https://www.w3.org/ns/ssn/.

[6] foaf: http://xmlns.com/foaf/0.1/.

[7] acl: http://www.w3.org/ns/auth/acl#.

[8] wo: http://smiy.sourceforge.net/wo/spec/weightingontology.html#Weight.

[9] void: http://vocab.deri.ie/void#Dataset.

[10] xsd: http://www.w3.org/2001/XMLSchema#.

Please cite this article as: O.R. Sanchez, I. Torre and B.P. Knijnenburg, Semantic-based privacy settings negotiation and management, Future Generation Computer Systems (2019), https://doi.org/10.1016/j.future.2019.10.024.

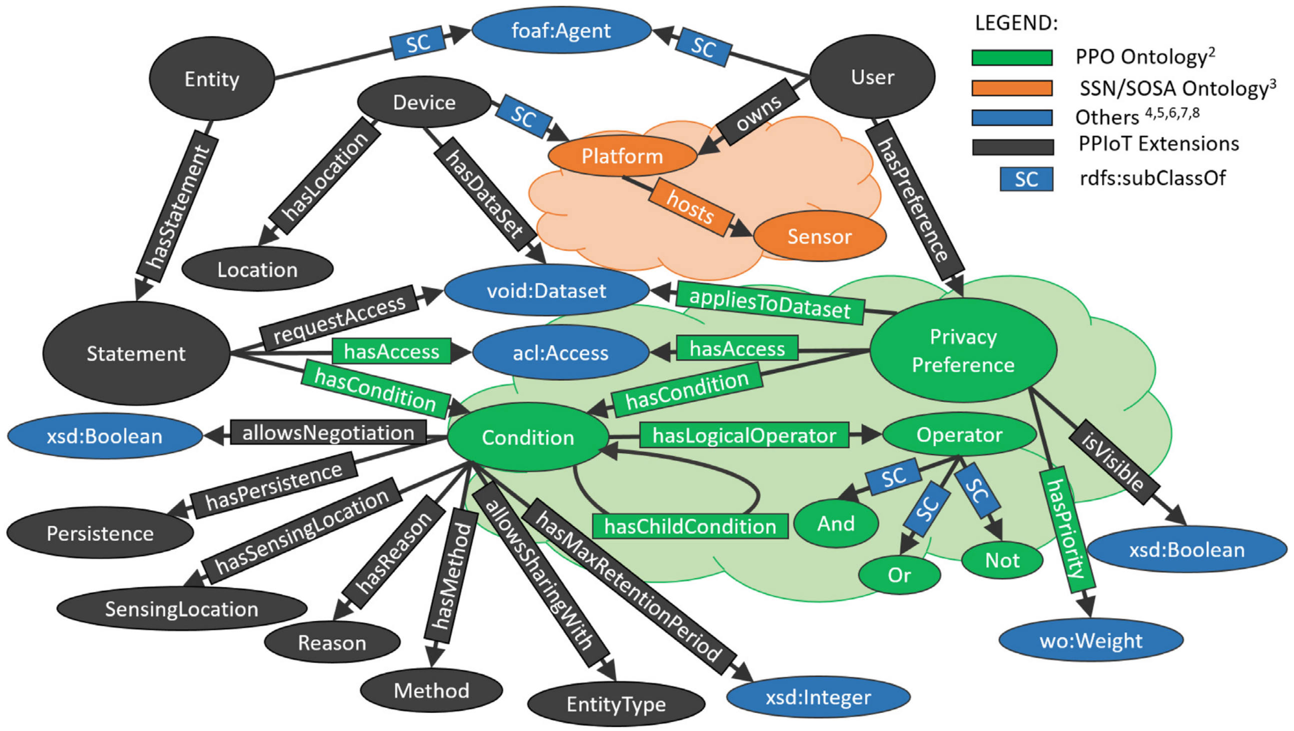

Fig. 1. Streamlined representation of the proposed Privacy Preference for IoT (PPIoT) Ontology. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

- ppo:hasCondition: the conditions of a user's privacy preference (that we extend to the new Class Statement, described below, so that hasCondition applies to Privacy Preference *or* Statement);
- ppo:hasLogicalOperator: the type of logical operator of the condition/child condition;
- ppo:hasChildCondition: used to create logical nested conditions in combination with the logical operators;
- ppo:hasAccess: the access control privilege which is granted by the user, described using the WAC vocabulary (we extend the property to the new Class Statement, described below, so that hasAccess applies to Privacy Preference *or* Statement)
- ppo:appliesToDataset: a privacy preference that applies to a Dataset instance;
- ppo:hasPriority: a value that signifies the rank of a privacy preference;
- sosa:hosts: the relation between the platform and sensor(s).

3.2. Extended classes and properties

Perera et al. argue that the complexity of the IoT paradigm demands that privacy approaches must offer more than the traditional allow or deny option and instead have room for negotiation between the user and TPs ('Entity' in the ontology) [24]. Furthermore, the extended classes and properties must at least be able to represent the conditions and the access conditions for both the Statement (for a TP) and Privacy Preference (for the user), the purpose/reason of collection, the persistence of access, the location, the retention period, and the usage method [16–19]. These principles are also included in the GDPR and FIP. We also included the common data-sharing schemes of TPs where they ask for permission to let other TPs access the user's data [50] and group them according to Entity type. Below, the major new classes and properties that implement the mentioned principles are briefly described. The mappings with *GDPR concepts* are provided in Section 2.4.2.

The new Classes are:

- User: the owner of the privacy preferences; a subclass of foaf:Agent;
- Entity: any agent that wants to access user information such as a human or a TP application—also a subclass of foaf:Agent but disjoint from User;
- Device: the specific IoT device of the User; a subclass of sosa:Platform;
- Location: the current location of the Device;
- Reason: the purpose of an Entity for accessing the User's data (e.g., health, social, fitness, etc.);
- Persistence: the frequency of data acquisition by the Entity;
- Method: how the data will be processed/utilized;
- SensingLocation: the location of an observation;
- EntityType: the type of Entity (for grouping purposes);
- Statement: the privacy policy Statement declaration of an Entity that consists of conditions regarding the request to access the user's dataset.

The new Properties are:

- owns: the relation between the User and her/his Device;
- hasPreference: a privacy preference of the User;
- hasLocation: the Location of the Device;
- hasDataset: the Dataset of the Device;
- hasReason: the Reason of the Condition;
- hasMethod: the Method of the Condition;
- hasPersistence: the Persistence of the Condition;
- hasSensingLocation: the SensingLocation of the Condition;
- allowsNegotiation: a boolean data type property that specifies whether the condition (set by the User or Entity) is negotiable;

Please cite this article as: O.R. Sanchez, I. Torre and B.P. Knijnenburg, Semantic-based privacy settings negotiation and management, Future Generation Computer Systems (2019), https://doi.org/10.1016/j.future.2019.10.024.

- hasMaxRetentionPeriod: An integer data type property that specifies the maximum retention period in hours of the data accessed by an Entity;
- hasStatement: the privacy Statement declared by an Entity;
- requestAccess: the Dataset(s) that the Entity requests in the Statement;
- allowsSharingWith: which type of Entity is allowed to share the accessed dataset;
- isVisible: a boolean data type property that specifies if the privacy preference of the User is visible to an Entity.

It is worth noting that both the user and the TP can *set which conditions are negotiable* for the PDM to optimize the negotiation process and recommendation. This can be expressed through the *allowsNegotiation* condition. Otherwise, the setting would be non-negotiable, as is the case in existing ontologies.

3.3. Ontology engineering and validation

The PPIoT ontology has been developed following the guidelines and steps stated in Noy and McGuinness [51]. This section discusses the development and validation of the proposed PPIoT ontology.

3.3.1. Domain modeling and ontology definition

Following the guidelines in [51], the first step in developing our ontology was the definition of the domain and scope. While our main goal was clear (i.e., representing privacy preferences from the perspective of both the user and the IoT TP, taking the GDPR requirements into account), the definition of the domain model was a non-linear, iterative process. This process started with the collection of relevant domain knowledge and the identification of use cases and competency questions, which were eventually used for the evaluation of our ontology.

In the spirit of the Linked Data paradigm, and following the principle of ontology reuse [51,52], we first analyzed existing ontologies that variously model concepts and relations in our domains of interest, namely the IoT domain and the privacy domain. For the IoT domain, we consulted the Linked Open Vocabularies for Internet of Things catalog (LOV4IoT[11]), which includes 510 ontology-based research projects in different IoT domains. A recent study [53] showed that, among the IoT ontologies, the W3C Semantic Sensor Network (SSN) ontology is the most commonly re-used ontology in other ontologies and can be considered as a de-facto IoT standard ontology. The SSN (which we selected as our *IoT*-related base ontology) also provides alignments to a variety of related ontologies and specifications.

With regard to the privacy domain, we analyzed the privacy research literature and the Linked Open Vocabularies (LOV[12]), which includes stable, high-quality ontologies. In addition, we consulted experts in privacy legislation. It is worth noting that at the time of designing the PPIoT ontology, while we found 8 vocabularies modeling privacy-related concepts (including the PPO ontology that we selected as our privacy-related base ontology), we did not find any GDPR-based vocabularies, since drafts of these were published after we developed our ontology (see Related Work in Section 2.4).

The PPO ontology was selected since it already models user privacy preferences for linked data, with a particular focus on social networks. Concepts and relations modeled therein fitted well within privacy modeling requirements in the IoT domain, even though extensions were needed.

The important terms of the domain were then enumerated [51]. Starting from the terms in the PPO ontology, we identified the missing terms for the IoT domain with respect to our goal. Our approach was to extend PPO by defining alignments with other ontologies and by introducing new terms when none were available in other ontologies. The principles related to personal data protection are explicitly stated in Article 5 of the GDPR Regulation [14]. These principles also guided the formulation of our *competency questions* in terms of the capability of the ontology to represent, and thus identify, the user privacy preference conditions for different datasets in terms of: lawfulness, fairness and transparency, purpose limitation, data minimization, accuracy, storage limitation, integrity and confidentiality (see Section 2.4). PPIoT is designed to include classes and properties that address the GDPR requirements for the management of personal data (see in particular the properties *hasReason*, *hasMethod*, *hasPersistence*, *hasSensingLocation*, *allowsNegotiation*, *hasMaxRetentionPeriod*, and *allowsSharingWith*).

The open-source ontology editor and framework Protégé[13] [54] was used to build the PPIoT ontology.

3.3.2. Ontology validation

We evaluated the PPIoT ontology using three methods [51,55]: (i) a coherence and consistency check, (ii) a task-based and application-based evaluation, and finally (iii) an evaluation using Competency Questions.

(i) Coherence and consistency check. The studies in [56,57] state that consistency validation refers to checking whether it is possible to obtain contradictory conclusions from valid input definitions: an ontology is logically consistent when it involves no logical contradiction. The PPIoT ontology was evaluated using different Reasoners in Protege. Reasoners provide consistency checks on the ontology, verifying that the ontology is logically consistent.

In our PPIoT ontology, none of the classes and axioms had logical contradictions. Fig. 2(a) shows that our proposed ontology is proven to be coherent and consistent using several Protégé reasoners (i.e., FaCT [58], HermiT [59], and Pellet [60]). A total of 50 axioms that are used in the PPIoT ontology were tested. We evaluated different reasoners since each reasoner performs differently for each task. In our case, all reasoners concluded that our ontology did not have inaccuracies. Fig. 2(b) shows the inferences and results using the Hermit reasoner, together with its computation time.

(ii) Task-based and application-based validation. According to [55,61], this type of validation involves evaluating how effective an ontology is in the context of a task or an application. In this light, the "application" may be an actual software program or a use-case scenario [55].

Application-based evaluation has been used, for example, to validate the PPO Ontology [62]. Sacco and Passant validated the ontology by building a privacy manager that could implement the creation of privacy preferences for RDF data described using PPO, and that could filter requested data by applying the preferences. [63] also proposed this method and measured the performance by comparing it to a gold standard.

In our case, we used our privacy manager (PDM) to perform the management of users' privacy preferences. The PDM prototype is available online.[14] The evaluation showed that the PPIoT ontology satisfies the requirements to model users' privacy preferences and TPs' request statements for user data. This prototype will be described in Section 4.3.

[11] https://lov4iot.appspot.com/
[12] https://lov.linkeddata.es/dataset/lov/
[13] https://protege.stanford.edu/
[14] https://github.com/OdnanOriginal/PDM

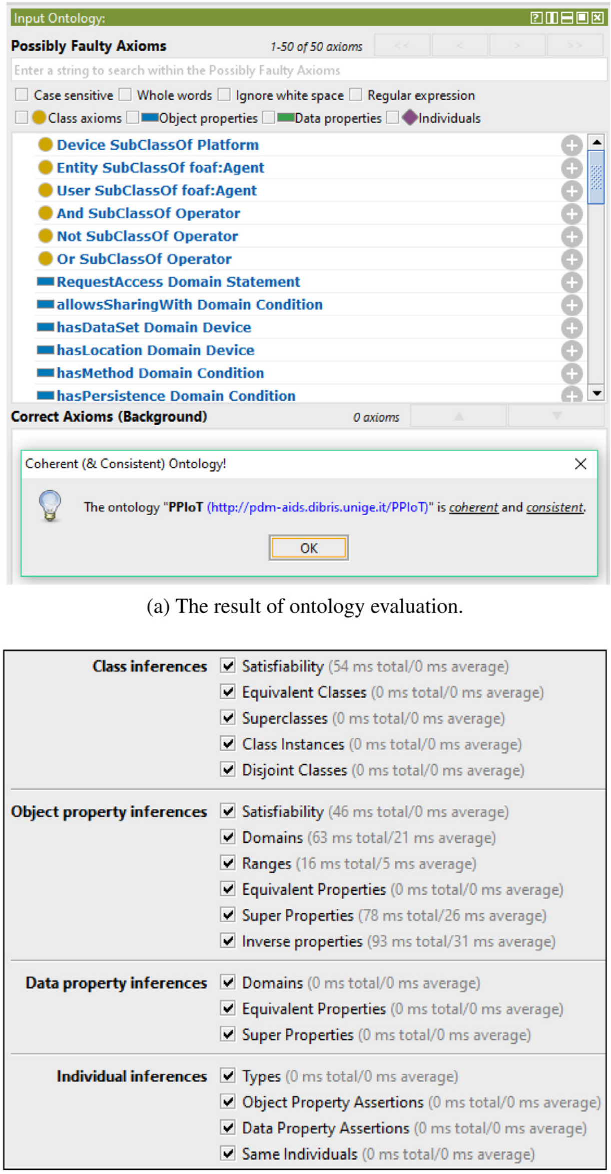

Input Ontology:

Possibly Faulty Axioms 1-50 of 50 axioms

Enter a string to search within the Possibly Faulty Axioms

□ Case sensitive □ Whole words □ Ignore white space □ Regular expression

□ ● Class axioms □ Object properties □ Data properties □ ◆ Individuals

- Device SubClassOf Platform
- Entity SubClassOf foaf:Agent
- User SubClassOf foaf:Agent
- And SubClassOf Operator
- Not SubClassOf Operator
- Or SubClassOf Operator
- RequestAccess Domain Statement
- allowsSharingWith Domain Condition
- hasDataSet Domain Device
- hasLocation Domain Device
- hasMethod Domain Condition
- hasPersistence Domain Condition

Correct Axioms (Background) 0 axioms

Coherent (& Consistent) Ontology!

The ontology "PPIoT (http://pdm-aids.dibris.unige.it/PPIoT)" is coherent and consistent.

OK

(a) The result of ontology evaluation.

Class inferences
☑ Satisfiability (54 ms total/0 ms average)
☑ Equivalent Classes (0 ms total/0 ms average)
☑ Superclasses (0 ms total/0 ms average)
☑ Class Instances (0 ms total/0 ms average)
☑ Disjoint Classes (0 ms total/0 ms average)

Object property inferences
☑ Satisfiability (46 ms total/0 ms average)
☑ Domains (63 ms total/21 ms average)
☑ Ranges (16 ms total/5 ms average)
☑ Equivalent Properties (0 ms total/0 ms average)
☑ Super Properties (78 ms total/26 ms average)
☑ Inverse properties (93 ms total/31 ms average)

Data property inferences
☑ Domains (0 ms total/0 ms average)
☑ Equivalent Properties (0 ms total/0 ms average)
☑ Super Properties (0 ms total/0 ms average)

Individual inferences
☑ Types (0 ms total/0 ms average)
☑ Object Property Assertions (0 ms total/0 ms average)
☑ Data Property Assertions (0 ms total/0 ms average)
☑ Same Individuals (0 ms total/0 ms average)

(b) The inferences tested by the HermiT reasoner in Protégé.

Fig. 2. The PPIoT ontology evaluation in Protégé.

In addition, we indirectly evaluated the PPIoT ontology by integrating it into our mock application. The mock application's understandability, control, simplicity and preferability were evaluated by real users on a 7-point Likert scale. The complete details of this evaluation are explained in Section 5.2.

While the aim of task/application-based evaluations is not to assess the generalizability of the ontology, but the performance of the ontology to support some tasks, generalizability could be addressed by applying this type of evaluation to more tasks in different applications.

(iii) Evaluation using Competency Questions. Competency questions can be used to design and then evaluate an ontology [51]. For example, this evaluation technique was used for the validation of the OSHCO ontology [64]: the authors developed competency questions for different use cases with the guidance of domain experts and then queried the ontology and checked the correctness of the retrieved answers. In our case, the identification of use cases and competency questions guided the design process, and they were used to validate the ontology as well. Our aim was to design an ontology that is able to answer questions about the management of user privacy preferences in IoT for different types of personal data and for different IoT domains. TP preferences for data requests made to users were modeled in the same way as user preferences. Thus, the two main *competency questions* that we aim to answer through the PPIoT ontology are the following:

- What are the user's privacy preference conditions for his/her different personal data (datasets, according to PPO and PPIoT terminology)?
- What are the TP privacy conditions required by TPs in their requests of personal data made to users?

More specific questions are aimed at identifying privacy conditions with respect to GDPR requirements (as explained in Section 3.2, from the side of the user and the TP), and to identify the user privacy preferences that are visible, i.e., can be queried by TPs. These questions can be answered by using SPARQL queries, as shown in Section 4.3 (Listings 3 and 4).

Ontology development is necessarily an iterative process and this process of iterative design will likely continue through the entire life-cycle of the ontology [51], in order to capture the domain changes and/or to align the ontology with new or updated ontologies that are modeled for that domain or sub-domains. We discussed this issue in the Related Works Section with regard to new GDPR-based ontologies that are being published, but this issue concerns the IoT domain as well.

3.4. PPIoT ontology running examples

This subsection shows how the PPIoT ontology can be used to set conditions for both the user privacy preferences and the TP statement.

3.4.1. User privacy preference

Listing 1 is an example of a privacy preference condition, *myPref*, in Turtle[15] notation. A user may have several conditions for different datasets. In this specific example, we present a user's privacy preference that applies to the *activity* dataset. Prefixes in Listing 1 are defined in Section 3.

The user preference has conditions which state that data access can happen only once (persistence), the maximum retention period of the data is 24 h, and the data is required to be encrypted if used for fitness-related reasons. These conditions are combined by the *LogicalOperator* ppo:And and can be negotiated (allowsNegotiation = true) except for the last condition (allowsNegotiation = false).

```
@prefix
    ppiot:<http://pdm aids.dibris.unige.it/PPIoT#>.
@prefix up:<http://www.userpreferenceExample.com#>.
up:userCond1 a ppo:Condition. up:userChildCond1 a
    ppo:Condition.
.....

up:myPref a ppo:PrivacyPreference;
  ppo:appliesToDataset ppiot:activity;
  ppo:hasCondition up:userCond1
    [ppo:hasLogicalOperator ppo:And;
    ppiot:hasPersistence ppiot:once;
    ppiot:hasMaxRetentionPeriod 24; #xsd:integer
```

[15] https://www.w3.org/TR/turtle/

```
     ppiot:allowsNegotiation true; #xsd:boolean
     ppo:hasChildCondition up:userChildCond1
       [ppiot:hasReason ppiot:fitness;
         ppiot:hasMethod ppiot:encrypted;
         ppo:hasLogicalOperator ppo:And;
         ppiot:allowsNegotiation false;
             #xsd:boolean;
         ppiot:allowsSharingWith
             ppiot:socialNetworkFriends];];
   ppo:hasAccess acl:Read, acl:Write;
   ppo:hasPriority wo:1;
   ppiot:isVisible true. #xsd:boolean
```

Listing 1: User preference condition

Notice that the last condition was expressed through a child condition with the *LogicalOperator* ppo:And (which is used as in [10]). This shows the significance of child conditions in letting users be more expressive with their privacy preference conditions. Moreover, *myPref* has "read" and "write" access permission, it has the maximum priority (value = 1) that rules out other privacy preferences, it is visible, so that "enhanced" TPs that utilize the PPIoT ontology to adjust their parameters, i.e., to conform to the user's conditions (if they are negotiable for this dataset).

3.4.2. *Third party policy statement*

Listing 2 shows an example of a statement of an enhanced TP. It shows that the TP would like to request access to the user's *activity*, *sleep* and *heart rate* datasets. The TP requests access to this data for fitness purposes, and promises to store the data encrypted, for a maximum retention period of 24 h. These conditions are combined by the *LogicalOperator* ppo:And and are negotiable. However, the TP requires continuous access to this data (persistence), and it does not allow negotiation on this particular point. Finally, it only requests "read" access to this data. The automatic negotiation of privacy preferences based on the user's privacy preferences (Listing 1) and the TP's policies (Listing 2) will be described in Section 4.3.2.

```
@prefix
     ppiot:<http://pdm aids.dibris.unige.it/PPIoT#>.
@prefix tpb:<http://www.TPExample.org#>.
tpb:cond1 a ppo:Condition. tpb:childCond1 a
     ppo:Condition.
.....
tpb:statementB a ppiot:Statement;
   ppiot:requestAccess ppiot:activity, ppiot:sleep,
        ppiot:heartRate;
   ppo:hasCondition tpb:cond1
        [ppo:hasLogicalOperator ppo:And;
          ppiot:hasReason ppiot:fitness;
          ppiot:hasMethod ppiot:encrypted;
          ppiot:hasMaxRetentionPeriod 24;
               #xsd:integer
          ppiot:allowsNegotiation true; #xsd:boolean
          ppo:hasChildCondition tpb:childCond1
               [ppiot:hasPersistence ppiot:continuous;
                 ppiot:allowsNegotiation false;
                      #xsd:boolean];];
   ppo:hasAccess acl:Read.
```

Listing 2: TP statement

4. Framework for privacy preference management

4.1. *Privacy preference model for interactive privacy setting*

Now that we have described the PPIoT ontology, we present our approach to support the user and the TPs to manage privacy preferences using our PPIoT-based Privacy Preference Model (PPM), shown in Fig. 3. The PPM model is composed of:

- the PPIoT-based Data Model, which formally specifies the privacy conditions that have to be taken into account when managing and processing personal data. Basically, this addresses the 'data management and processing requirements' described in Section 2.4 and reported in Fig. 3,
- the Interaction Model, which addresses the 'communication and transparency' requirements. This model consists of an interactive approach that requests users' explicit consent to each condition specified in the TP statement based on the PPM data model.

Fig. 3 shows that the Interaction model can be instantiated in different ways, depending on whether the user has a Personal Data Manager (PDM in the figure) that mediates the interaction with the IoT TP and depending on the TP type. If a PDM is not available, the Interaction model can be instantiated as an *Interactive User Interface*. Otherwise, it would be instantiated as an *Interactive Negotiation and Recommendation process* managed by a Personal Data Manager. The former approach, "PPM-based Interactive User Interface" in the figure, can be adopted by TPs for direct interaction with the user. We have developed a mockup of this approach for a fitness application.[16] [65] This mockup improves upon traditional policy statements, which are usually presented in complex and lengthy "terms and conditions" that users rarely actually read [13,14]. As shown in the figure, the PPM-based Interactive UI conforms to the PPM data model – and consequently the PPIoT ontology – and is aimed at supplementing the traditional privacy policy. The interaction design and layout of the mockup follow the GDPR requirements of transparency, simplicity, and explicit consent for each request. The TP can store the user's consent data in an RDF store or any other database and can autonomously define its strategies for recording the origin and use of data related to such consent, as required by GDPR, and for tracking its changes over time (in this respect, see for example the GDPRprov ontology [49] for provenance modeling, discussed in the Related Work Section).

With reference to Fig. 3, if a PDM is available, it can be used to mediate the interaction and privacy setting process between the user and the TP. In this case, the PPM is used by the PDM to conduct an interactive negotiation and recommendation of privacy preferences. The goal, in this case, is to further simplify the management of privacy preferences, relieving the user from the burden of specifying her/his preference conditions for each new device and application, but maintaining a certain level of control. The focus of this paper is on the interactive settings through the PDM (which will be described in the remainder of the paper), but it is worth noting that the interactive user interface

[16] http://pdm-aids.dibris.unige.it/simulation.php

Please cite this article as: O.R. Sanchez, I. Torre and B.P. Knijnenburg, Semantic-based privacy settings negotiation and management, Future Generation Computer Systems (2019), https://doi.org/10.1016/j.future.2019.10.024.

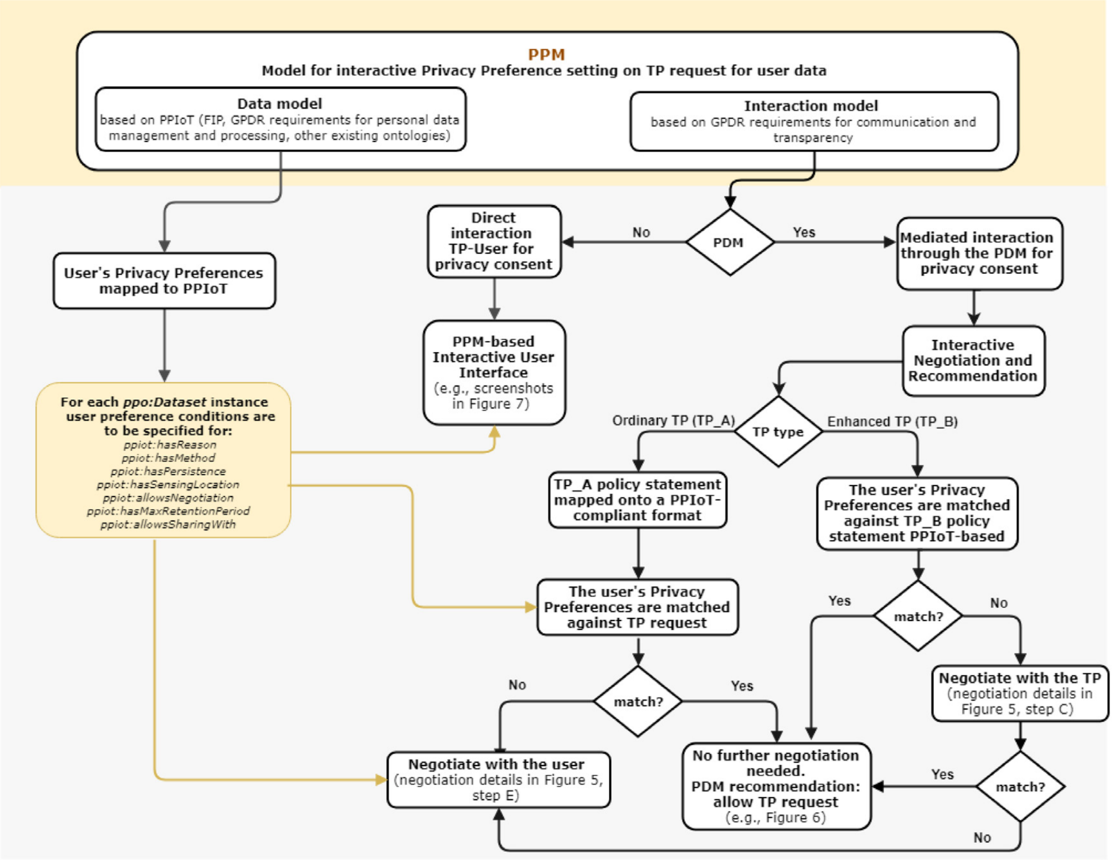
PPM
Model for interactive Privacy Preference setting on TP request for user data
Data model
based on PPIoT (FIP, GPDR requirements for personal data management and processing, other existing ontologies)
Interaction model
based on GPDR requirements for communication and transparency
User's Privacy Preferences mapped to PPIoT
For each *ppo:Dataset* instance user preference conditions are to be specified for:
ppiot:hasReason
ppiot:hasMethod
ppiot:hasPersistence
ppiot:hasSensingLocation
ppiot:allowsNegotiation
ppiot:hasMaxRetentionPeriod
ppiot:allowsSharingWith
PDM
No
Yes
Direct interaction TP-User for privacy consent
Mediated interaction through the PDM for privacy consent
PPM-based Interactive User Interface (e.g., screenshots in Figure 7)
Interactive Negotiation and Recommendation
TP type
Ordinary TP (TP_A)
Enhanced TP (TP_B)
TP_A policy statement mapped onto a PPIoT-compliant format
The user's Privacy Preferences are matched against TP_B policy statement PPIoT-based
The user's Privacy Preferences are matched against TP request
match?
Yes
No
Negotiate with the TP (negotiation details in Figure 5, step C)
match?
No
Yes
Negotiate with the user (negotiation details in Figure 5, step E)
No further negotiation needed.
PDM recommendation: allow TP request (e.g., Figure 6)
match?
Yes
No

Fig. 3. Privacy Preference Model (PPM) for interactive privacy setting.

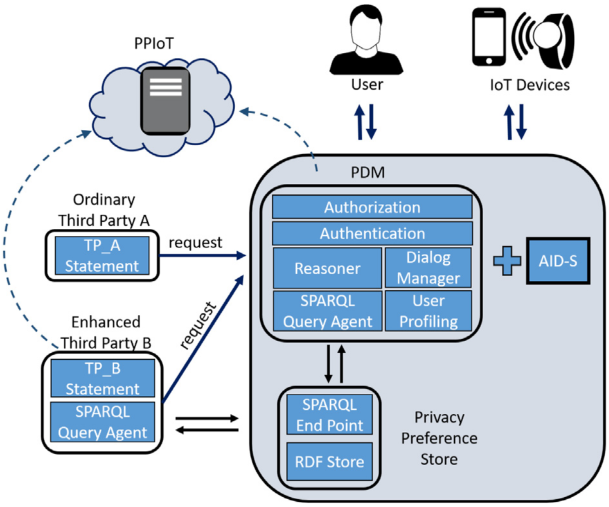

Fig. 4. Overview of the framework implementing the Privacy Preference Model.

is based on the same PPIoT-based PPM. The right part of Fig. 3 depicts the workflow handled by the PDM. The Interaction model is designed to take into account both ordinary TPs (TP_A), which do not adhere to this framework and are therefore not aware of the PDM mediation, and enhanced TPs (TP_B), which define their Policy statement in accordance with the PPIoT and negotiate with the PDM based on the PPM. The details of the negotiation in Fig. 3 will be explained in the next sections.

4.2. Interactive negotiation and recommendation of privacy preferences through a personal data manager

We now turn to the use of our PPIoT ontology in the most technologically advanced scenario, which is the PPM-based interactive negotiation and recommendation through a Personal Data Manager (PDM) and an Adaptive Inference Discovery Service (AID-S). This framework, shown in Fig. 4, was first defined in Torre et al. [50,66]. The PDM is responsible for managing the user's privacy preferences and the interaction with the TPs. It also has the tasks of authorization, authentication and user profiling. The AID-S is responsible for the computation of potential inference risks given the combination of personal data requested by the TP, and for giving users suitable recommendations regarding these risks. Further details on AID-S for the specific case of fitness tracking can be found in Torre et al. [50] and will not be discussed further in this paper.

In this paper, we aim to present the PDM's adoption of the SWT for the management of the user's privacy preferences and the interaction between the user and the TPs. Fig. 4 shows how the PDM acts as an intermediary between the user, her/his personal IoT devices, and the TPs. The user's privacy preferences, annotated with the concepts of the PPIoT ontology, are stored in an RDF store and made available to the PDM Query Agent through a SPARQL endpoint. The PDM can be implemented as a client–server application, as a service in the cloud, or even as a semantic mobile app with a mobile endpoint (for an example implementation of this solution see Yus and Mena [67]). The current implementation is a client–server application where the client runs as a mobile app while the RDF store and the reasoner are on the server.

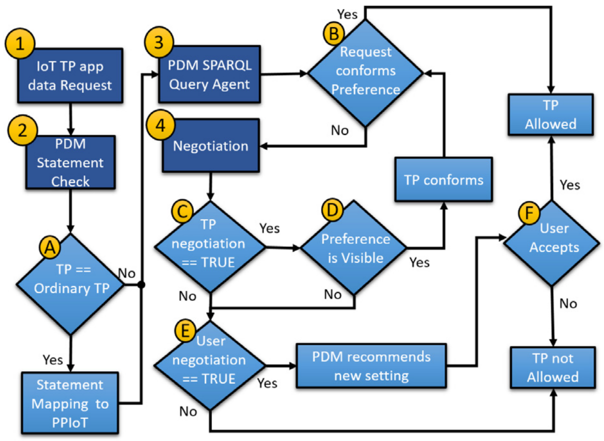

Fig. 5. The simplified interaction workflow between the PDM, the TP and the user.

The TPs shown in Fig. 4 include an ordinary TP (TP_A) and an enhanced TP (TP_B) that has SWT capabilities and utilizes the PPIoT ontology. For TP_A, the dialog manager maps its policy statement onto a PPIoT-compliant format. Since the ordinary TP has no capabilities to negotiate, the PDM cannot facilitate any negotiation between the user and the TP, so the "negotiation" (i.e., whether to accept the policy or not) will only be between the user and the PDM, as we will describe in Section 4.3.3. TP_B, on the other hand, will be able to benefit from the negotiation capabilities embedded in the PPIoT ontology, effectively giving the TP and its users the option to specify individual privacy settings, conditions on these settings, and the negotiability of such conditions.

4.3. Use Case demonstration for privacy settings negotiation and recommendation

In this section, we will show how the PDM uses the PPIoT ontology and how it performs the negotiation between the user and the TP. The PDM prototype[17] was developed using the Java programming language, the Jena Semantic Web Framework[18] and an Apache Jena Fuseki SPARQL server for RDF storage and querying the user's preference data. The open-source ontology editor and framework Protégé[19] was used to build the PPIoT ontology.

This use case answers the *competency question* "what are the user's privacy preference conditions for different datasets". It also provides a consistency check on the ontology, verifying that the ontology is logically consistent through the Jena reasoner. None of the classes and axioms had logical contradictions. Fig. 5 shows a simplified version of the PDM workflow. Its four main steps provide the core phases of negotiation, which will be elaborated below.

4.3.1. TP application data request and PDM statement check

In step 1 (Fig. 5), a TP Statement will be issued during the installation or update of an application. In this instance, the PDM acts as a dialog manager and mediates the interaction between the user and the TP. Our instantiation of the PDM as a mobile application is designed to have the capability to interrupt the installation and check the permissions requested by the TP.

Step 2 is the interpretation of the Statement. In this step, decision block A checks if the requesting TP is an ordinary TP or an enhanced TP. For an ordinary TP, the PDM dialog manager can locate its Privacy Statement, which is usually stated in a file (e.g., Androidmanifest.xml[20] for Android apps), and map it onto the PPIoT ontology. An enhanced TP utilizes the PPIoT ontology to present its Statement, so no mapping is required.

4.3.2. PDM SPARQL query agent

Step 3 is the evaluation of the Statement. In this step, the PDM checks if the Statement conforms with the user's privacy preferences. The preferences are queried through the SPARQL Query Agent component. Listing 3 is an example of a query from the Query Agent to the privacy preference store that retrieves the list of user's privacy preferences for each dataset. It refers to the example presented in Section 3.4.

```
SELECT ?pref ?value
WHERE {?pref <ppo:appliesToDataset> ?value.}
```

Listing 3: A query of the user's privacy preferences for each dataset.

```
SELECT ?cond ?value
WHERE { <:myPref> <ppo:hasCondition> ?tempVariable.
        ?tempVariable ?cond ?value.}
```

Listing 4: A query example of conditions associated to a privacy preference.

By specifying a dataset (e.g., *activity*) for the *?value* variable, the Query Agent can retrieve all user preference conditions associated with this dataset. Subsequently, the PDM can retrieve all associated conditions and their values using the URI name of the privacy preference as input (Listing 4). Fig. 6(a) shows the result of this PDM query. The query for ppiot:hasAccess, ppo:hasPriority, and ppo:haschildCondition can be done in a similar manner.

If the TP's request conforms with the user's privacy preferences, it will pass the statement check (decision box B == false) and be granted access. The PDM can now act as an intermediary for the disclosure of the information that is included on the conditions of the user (Privacy Preference) and the TP (Policy Statement) specified in the prior steps.

4.3.3. Negotiation with an ordinary TP

If the TP's request does not conform with the user's preferences (decision box B == false), negotiation is needed. There are two cases for this negotiation. If the TP is an ordinary TP (i.e., TP_A), there is no opportunity for negotiation on the TP's side (decision box C == false), and negotiation will only be possible between the user and the PDM if the user has indicated her/his preferences as negotiable (decision box E == true). In this case, the PDM will provide a recommendation to allow the negotiable preference, which the user can accept or deny (decision box F).

[17] https://github.com/OdnanOriginal/PDM

[18] https://jena.apache.org/index.html

[19] https://protege.stanford.edu/

[20] https://developer.android.com/guide/topics/manifest/manifest-intro.html

Please cite this article as: O.R. Sanchez, I. Torre and B.P. Knijnenburg, Semantic-based privacy settings negotiation and management, Future Generation Computer Systems (2019), https://doi.org/10.1016/j.future.2019.10.024.

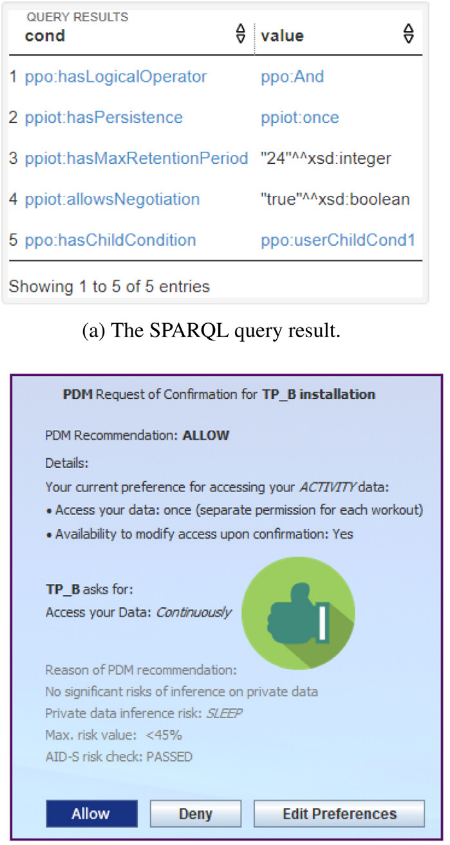

QUERY RESULTS

cond	value
1 ppo:hasLogicalOperator	ppo:And
2 ppiot:hasPersistence	ppiot:once
3 ppiot:hasMaxRetentionPeriod	"24"^^xsd:integer
4 ppiot:allowsNegotiation	"true"^^xsd:boolean
5 ppo:hasChildCondition	ppo:userChildCond1

Showing 1 to 5 of 5 entries

(a) The SPARQL query result.

PDM Request of Confirmation for **TP_B installation**

PDM Recommendation: **ALLOW**

Details:

Your current preference for accessing your *ACTIVITY* data:

- Access your data: once (separate permission for each workout)
- Availability to modify access upon confirmation: Yes

TP_B asks for:

Access your Data: *Continuously*

Reason of PDM recommendation:

No significant risks of inference on private data

Private data inference risk: *SLEEP*

Max. risk value: <45%

AID-S risk check: PASSED

(b) The PDM confirmation request and recommendation.

Fig. 6. The query result of Listing 4 and the PDM recommendation to the user.

Example. Considering the conditions in the example described in Listings 1 and 2, the PDM finds that the conditions for the TP request to access the activity dataset comply with the user privacy preference regarding the *reason*, *method*, and *maxRetentionPeriod*. However, the *persistence* condition requested by the TP is "continuous" while the user preference is set to "once" for the activity dataset. The PDM finds that the user allows negotiation (decision box E == true), so it recommends giving TP_A "continuous" access to the *activity* dataset (see Fig. 6(b)). If the user does not allow negotiation (decision box E == false), the TP's request will be denied, that is, the user will be recommended not to allow.

4.3.4. Negotiation with an enhanced TP

In the case of an enhanced TP (e.g. TP_B), both the user and the TP can set negotiation values. The PDM will then first check if the TP allows negotiation on this aspect (decision box C == true) and if the user's preference for this aspect is set to "visible" (decision box D == true). If so, the enhanced TP can query the user preference using Listings 3 and 4.[21] and modify its request in order to conform with the preferences (decision box B == true). If either the TP does not allow negotiation (decision box C == false) or the user's preference is not set to "visible" (decision box D == false), then negotiation will continue between the user and the PDM alone (decision box E) as described above.

Example. Considering the conditions in the example described in Listings 1 and 2, upon finding the conflict regarding the *persistence* condition, the PDM will first check if the TP allows negotiation (decision box C) using the algorithm in Fig. 5. Checking the TP for negotiation first prioritizes the user's preference over the TP request. If the TP allows negotiation (decision box C == true), it will conform to the user's preferences, given that these preferences are visible (decision box D == true). An enhanced TP can query the user preference (Section 4.3.2 shows how to query the preference store using SPARQL queries) if it is set to visible by the user through the *isVisible* property. It can then conform by either removing the request of those data on its Statement that have conditions that conflict with the user's preference or by changing these conditions in accordance to the user's preference. Unfortunately, the enhanced TP_B in the example (Listing 2) does not allow negotiation (decision box C == false). This could for instance happen if the TP is a fitness tracker, which needs continuous access to activity data to keep track of the user's calories burned. Accordingly, the PDM then checks the user negotiation conditions in Listing 1. The PDM finds that the user allows negotiation (decision box E == true). Therefore, the PDM proposes to the user to change the condition for *persistence* to "continuous" (decision box F). Fig. 6(b) shows the confirmation request to the user.

4.3.5. Evaluating inference risks

In the complete framework, the PDM also calls AID-S for the recommendation to detect whether the request generates any inference risk. In Fig. 6(b), AID-S computation of the inference risk for the datasets asked by the TP is considered to be low (which is shown in gray since it is not part of the paper). If the user thinks the inference risk is too high or anyway if she/he does not agree to grant the permission to the TP, the user can deny the request after all, or edit her/his privacy preferences. More details on the AID-S risk computation can be found in Torre et al. [50]. It is worth noting that PDM recommendation to allow or to reject a TP request works also without the inference risks check, which is an add-on to support the user choice. In any case, the result of the process is a recommendation since the user has always to provide explicit consent to TP requests, in accordance to GDPR principles.

5. Evaluation

The offline evaluation of the automated negotiation process is done using an actual running application (cf. Section 4.3), where a use case was simulated and explained step by step. Automation is hard to evaluate in a test environment, though, since it requires simulating all the interaction conditions. Hence, we focus here on evaluating the fully interactive version of PPM, noting that the underlying PPIoT ontology is the same. In this section we describe the results of our online study with users evaluating the impact of the PPM.

As explained in Section 4.1, the goal of the PPM is to provide an interactive approach that enables TPs to use our PPIoT ontology vocabularies to improve control and transparency in the presentation of privacy policy statements. Section 4.3 demonstrated how the Personal Data Manager can use the PPM for negotiating and interactively setting the user's privacy preferences. Also, it shows the value and effectiveness of using the SWT by describing TP policy statements and users' preferences with the PPIoT ontology. As Fig. 3 describes, the alternative means to use the PPM is through an interactive user interface. In this section, we describe the results of a preliminary user evaluation of this user interface. Note that the use of the PPM through the PDM involves an interactive user interface as well, as shown in the Interaction Model in Fig. 3 —the main difference is that the PDM automates some of the interaction between the user and the TPs.

[21] This explains the request arrows from both the PDM and the enhanced TP_B to the Privacy Preference Store in Fig. 4

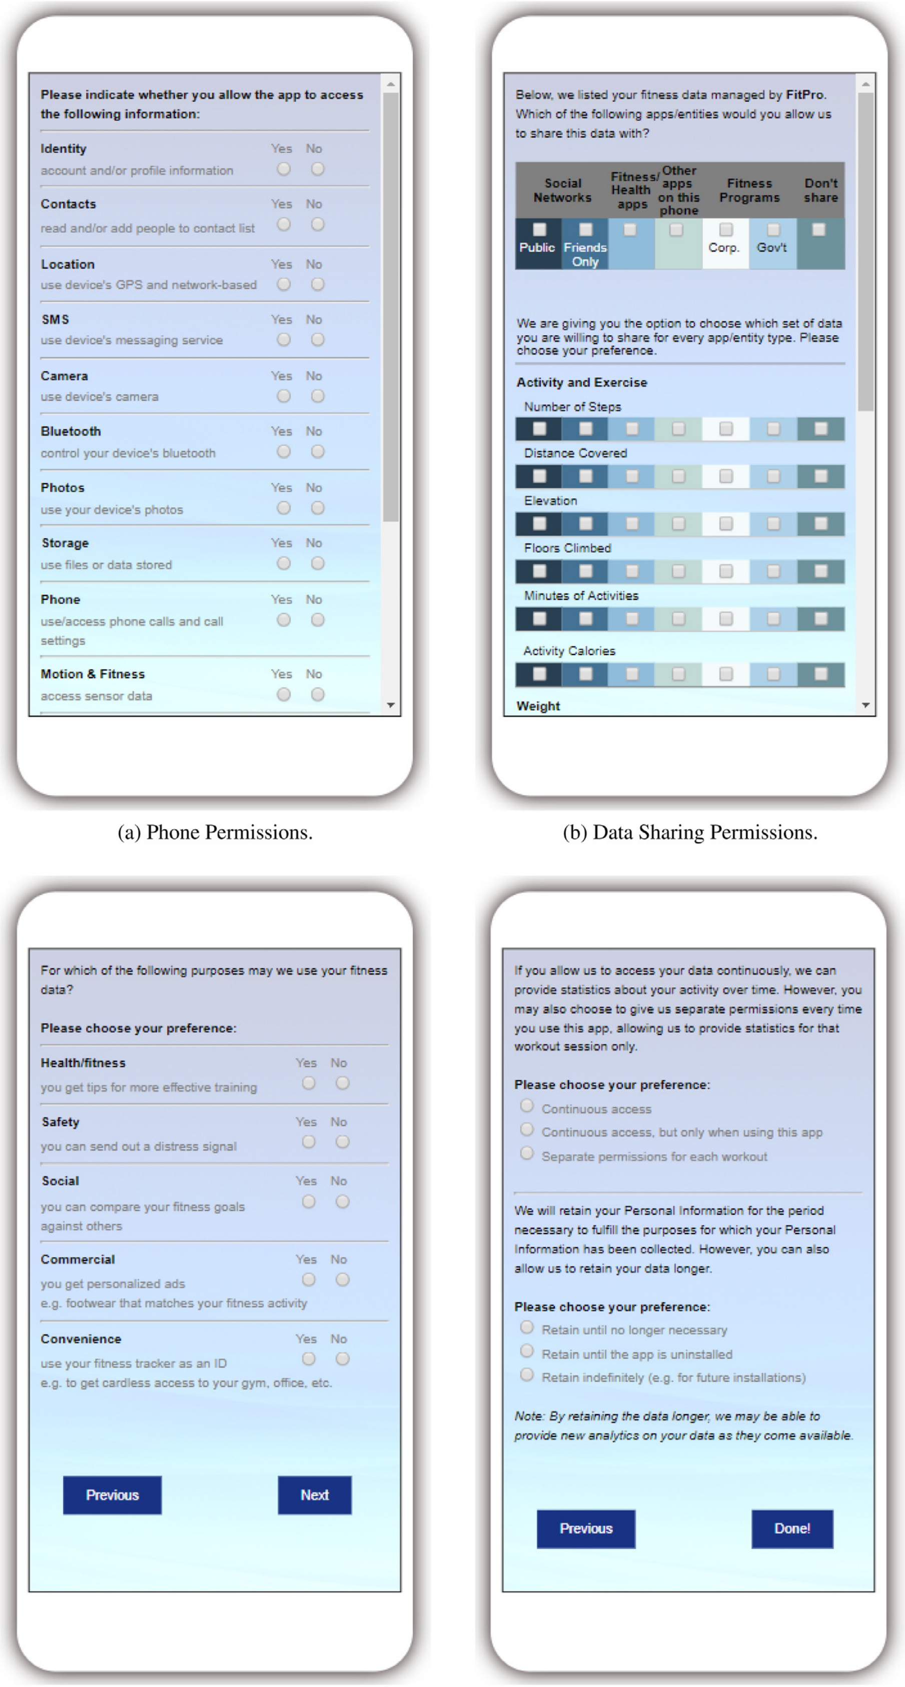

Please indicate whether you allow the app to access the following information:

	Yes	No
Identity account and/or profile information	○	○
Contacts read and/or add people to contact list	○	○
Location use device's GPS and network-based	○	○
SMS use device's messaging service	○	○
Camera use device's camera	○	○
Bluetooth control your device's bluetooth	○	○
Photos use your device's photos	○	○
Storage use files or data stored	○	○
Phone use/access phone calls and call settings	○	○
Motion & Fitness access sensor data	○	○

(a) Phone Permissions.

Below, we listed your fitness data managed by **FitPro**. Which of the following apps/entities would you allow us to share this data with?

	Social Networks		Fitness/Health apps	Other apps on this phone	Fitness Programs		Don't share
	□ Public	□ Friends Only	□	□	□ Corp.	□ Gov't	□

We are giving you the option to choose which set of data you are willing to share for every app/entity type. Please choose your preference.

Activity and Exercise

Number of Steps	□	□	□	□	□	□	□
Distance Covered	□	□	□	□	□	□	□
Elevation	□	□	□	□	□	□	□
Floors Climbed	□	□	□	□	□	□	□
Minutes of Activities	□	□	□	□	□	□	□
Activity Calories	□	□	□	□	□	□	□

Weight

(b) Data Sharing Permissions.

For which of the following purposes may we use your fitness data?

Please choose your preference:

	Yes	No
Health/fitness you get tips for more effective training	○	○
Safety you can send out a distress signal	○	○
Social you can compare your fitness goals against others	○	○
Commercial you get personalized ads e.g. footwear that matches your fitness activity	○	○
Convenience use your fitness tracker as an ID e.g. to get cardless access to your gym, office, etc.	○	○

(c) Purpose of Collection.

If you allow us to access your data continuously, we can provide statistics about your activity over time. However, you may also choose to give us separate permissions every time you use this app, allowing us to provide statistics for that workout session only.

Please choose your preference:

○ Continuous access
○ Continuous access, but only when using this app
○ Separate permissions for each workout

We will retain your Personal Information for the period necessary to fulfill the purposes for which your Personal Information has been collected. However, you can also allow us to retain your data longer.

Please choose your preference:

○ Retain until no longer necessary
○ Retain until the app is uninstalled
○ Retain indefinitely (e.g. for future installations)

Note: By retaining the data longer, we may be able to provide new analytics on your data as they come available.

(d) Frequency & Retention Period.

Fig. 7. PPM-based Interactive UI for Privacy Policy Settings conforming the PPIoT ontology.

5.1. Interactive user interface for the privacy preference model

For the evaluation of the PPM-based Interactive User Interface (UI), we used the mockup of a fitness application called "FitPro"[22] which was introduced in Section 4.1. FitPro presents an interactive PPIoT-based privacy policy to its users. Below we discuss the permissions requested by the mockup, which were taken from existing fitness trackers that have the capability to share fitness data (i.e., Fitbit,[23] Misfit,[24] Jawbone,[25] Garmin[26]). To make our study more generalizable to the fitness domain, our mockup requests the superset of all phone and fitness data permissions of these fitness trackers.

Fig. 7 shows the user interface of the *FitPro* set-up pages that query users about their privacy preferences. Figs. 7(a) and 7(b) request smartphone and fitness data permissions, respectively. We use the ask-on-installation (AOI) permission model [68], as it simplifies our study design. The ask on first use (AOF) model, which is currently used by Android 6.0+ and iOS 5.0+, also works with the PPIoT-based PPM, since it conceptually does not matter when the app asks the user for permissions.

The fitness data page in Fig. 7(b) asks the user with whom (i.e., which other TPs) the current TP (FitPro) is allowed to share access to her/his Fitness data. Such data sharing is not uncommon in fitness apps, which allow users to enjoy a plethora of services that other TPs have to offer in addition to the services from the main TP. As such, this page refers to the *allowsSharingWith* property condition and the *Dataset* class, that is, it addresses both the *who* and *what* dimensions of the privacy settings simultaneously, allowing users to set different permissions for different types of TPs from this central interface. Most current applications use a "federation" approach, where the user is directed to the main TP every time another TP requests access to the user's fitness data. This becomes a tedious task for the user, especially if he/she has several other TP apps. The PPIoT-based PPM could help reduce this redundancy and gives users more centralized control [7].

Fig. 7(c) shows the page that asks users the permitted *purposes* of data collection (which refers to the *hasReason* property), while Fig. 7(d) asks users about the permitted *frequency* of collection (*hasPersistence* property) and the maximum *retention* period for the accessed data (*hasMaxRetentionPeriod* property). For this particular example, we did not consider the *method* of data collection (encrypted or unencrypted), but a TP may interactively inquire about permitted data protection methods as well.

5.2. Sample and methodology

For the evaluation, we recruited a total of 310 Fitbit fitness tracker users via the Amazon Mechanical Turk[27] crowd-sourcing platform linked to our test environment.[28] We restricted participation to fitness tracker users to be able to compare the current privacy settings of their existing fitness app against their settings in our FitPro mockup. Moreover, we restricted participation to Fitbit users to reduce the app-based variability in privacy settings among our participants (cf. each fitness tracker requests slightly different permissions and personal information from its users). Note, though, that our FitPro mockup contains permission requests from a variety of fitness trackers, and hence our results generalize beyond Fitbit to fitness trackers in general.

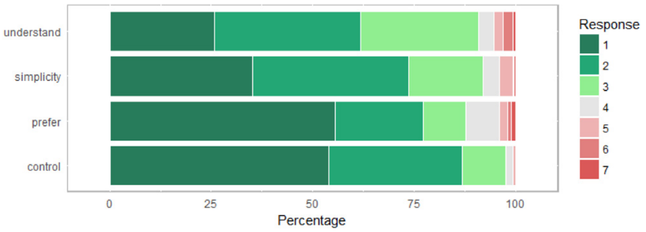

Fig. 8. 7-point scale evaluation on the PPIoT-based PPM.

We removed data from 15 participants whose completion times and answers to the control questions clearly indicated a lack of attention to the study, resulting in a final dataset of 295 responses. The participants are composed of 34.2% males (101 participants) and 65.8% females (194 participants), had a mean age of 35, and were generally highly educated (62% had at least a bachelor's degree).

Participants were asked to use the *FitPro* user interface as if they were installing a new application on their device. The participants were subsequently asked to respond to a questionnaire asking for their feedback about this installation experience, their current Fitbit settings, and their privacy preferences.

5.3. Subjective evaluation

The participants were asked (using 7-point scale items) about the *understandability* of the presented privacy policy, the amount of *control* they thought the interface gave them, how *easy or difficult* is was to set their privacy preferences, and whether they *preferred* this PPM-based interactive settings interface over the traditional privacy policy statements they experienced when installing their apps. Participants' feedback is presented in Fig. 8, and the following are the averages for each item (lower = better) combined for all respondents:

- Understandability = 2.31 (1 = Definitely Understand, 7 = Definitely do not understand)
- Control = 1.62 (1 = Definitely gave control, 7 = Definitely did not give control)
- Simplicity = 2.04 (1 = Very easy, 7 = Very difficult to use)
- Preferability = 1.86 (1 = Definitely Prefer, 7 = Definitely do not prefer over the traditional privacy preference model)

The results of our evaluation show that the interactive PPM-based interface helps participants understand the TP's privacy policy, making clear the options and presenting them as a structured, interactive dialog. Participants tend to prefer this PPM-based interactive privacy policy over a traditional privacy statement, which is unsurprising given that so few users actually read such statements [13,14]. Arguably, the interactive PPM format is easier for users to engage with, comprehend, and retain. This is reflected in the fact that 85% of survey respondents said they understood the privacy policy. This is a high number in light of the fact that many commercial privacy policies are notoriously hard to understand [69].

Overwhelmingly, the interactive PPM-based interface gives participants more control than a traditional presentation of TP's privacy policy and this is a consequence of the GDPR principles and the interactive presentation. Moreover, despite the complexity that comes with granular control, 90% of the participants find the interface easy to use. Admittedly, users would likely consider it a burden to make a large number of privacy decisions for a multitude of applications, and/or to frequently revisit these decisions as their privacy preferences evolve. This is where PDM can offer relief in the form of privacy recommendations.

[22] http://pdm-aids.dibris.unige.it/simulation.php
[23] https://www.fitbit.com/
[24] https://misfit.com/
[25] https://www.jawbone.com/
[26] https://www.garmin.com/
[27] https://www.mturk.com/
[28] http://pdm-aids.dibris.unige.it/index.php

Please cite this article as: O.R. Sanchez, I. Torre and B.P. Knijnenburg, Semantic-based privacy settings negotiation and management, Future Generation Computer Systems (2019), https://doi.org/10.1016/j.future.2019.10.024.

Overall, then, about 80% of the participants prefer the PPM-based interface over traditional privacy policies, with over 50% of participants definitely preferring it. Respondents are quite unanimous in their feedback, as standard deviations for these items are low (Understandability: 1.17, Simplicity: 1.25, Preferability: 0.8, Control: 1.07). We also find no significant differences in these evaluations (p-values > 0.05) in terms of gender, age or mobile OS. This shows that GDPR-based PPM does not only conforms to EU requirements for privacy but also results in a more appreciated approach to get consent (i.e., permissions on requested data).

5.4. The effectiveness of the PPM for the elicitation of privacy preferences

The above subjective evaluation suggests that our FitPro PPM-based interactive settings are an improvement over the privacy-setting experience of existing fitness apps. The rates at which various permissions were allowed by users in FitPro are displayed in Figs. 9 and 10. Permissions are grouped into the four sets requested in the FitPro simulated installation: In-app requests, Smart phone permissions, Fitness data (Fig. 9) and GDPR permissions (Fig. 10).

Our goal in this further evaluation is to compare such FitPro PPM-based permissions against participants' current permissions given to their *existing fitness apps*. In our users' case, the existing fitness apps are: (i) the participant's Fitbit app and (ii) the third-party apps on the participant's device that request permission to access their Fitness data managed by Fitbit.

It is worth recalling that FitPro simulates the installation of a fitness tracking app like Fitbit. In such installation it requests permissions on personal data (in-app requests), permission to access smartphone data (smartphone permissions), permissions to access, process and store such data (GDPR permissions). Moreover, it includes requests of permissions for sharing FitPro fitness data (Fitness data permissions) with other apps—which simulates third-party apps requesting Fitbit fitness data.

Thus, in principle, we can compare the settings that the participants set in the FitPro PPM-based system against those in their existing fitness apps. However, this comparison is not always possible or meaningful. Particularly, it is not possible when the participants do not have any current settings (i.e., no third-party app accessing their Fitness app). Moreover, it is not meaningful when Fitbit's settings are mandatory (i.e., mandatory app request permissions and mandatory allow-all blocks of permissions). In both cases, instead of using the participant's Fitbit settings, we asked participants about their preferences separately (questionnaire-reported preferences). For each set of permissions, except for the GDPR set (this data was not available on participants' devices or even in their experience since it is a novel contribution), we compared the PPM-based settings with the available settings as follows:

- For the 'in-app set', we compare participants' PPM-based settings against their self-reported[29] preferences to adhere to in-app data requests.
- For the 'smartphone set', we compare participants' PPM-based settings against their current settings, i.e., the actual permissions they have given to their Fitbit app (note that the requests differ slightly between Android and iOS).
- For the 'fitness data set', we have two situations: for users who have a third-party application accessing their current Fitbit data, we compare their PPM-based settings against the current settings for one of these third-party applications. For users who do not have any third-party applications, we compare their PPM-based settings against their self-reported preferences.

Table 1
Chi-square tests of association between PPM settings and participants' preference on in-app data requests.

In-app request (A set)	Preferences vs PPM settings
First name	10.6 ($p < 0.05$)
Last name	11.0 ($p < 0.05$)
Birth date	6.7 ($p < 0.05$)
Gender	1.3 ($p > 0.05$)
Height	0.2 ($p > 0.05$)
Weight	1.4 ($p > 0.05$)

We used chi-square tests to test the association between the settings mentioned above. Tables 1, 2, and 3 show the results of the chi-square tests, which will be explained in the following sections.

Large χ^2 values with small p-values (i.e., p-value < 0.05) indicate strong association. Overall, we found that participants' preferences expressed through the questionnaire are strongly associated with their PPM-based settings on FitPro, and the same happens for settings given when the user can freely allow or deny each permission, as in Fitbit phone permissions. Conversely, PPM-based settings are not significantly associated to the current settings of regarding the use of fitness data by other third-party apps.

The overall results seem to suggest that the more transparent and controllable PPM-based interactive setting, besides being more appreciated by participants, would also have an impact on the effectiveness in expressing their privacy preferences, even though further investigations are required to generalize these findings.

Below we discuss the results of these association tests for the sets requested in the FitPro simulated installation.

5.4.1. In-app request permissions

Fitness apps regularly ask users for their personal data such as name, surname, age, height and weight during sign-up. In most apps, this is compulsory information. In our study, however, we asked the participants if they would allow or deny such permissions if they were optional instead of required. From here, we compared participants' preferences with their PPM-based privacy settings. As depicted in Table 1, participants' preferences and PPM settings are strongly associated for first name, last name and birth date. Interestingly, these are also the most sensitive data in this group (see Fig. 9). For the other items, there is no significant association between participants' preferences and their FitPro PPM-based settings.

5.4.2. Smartphone permissions

Participants' average acceptance rates for smartphone permissions varies widely; participants were least likely to give FitPro access to their contacts and photos, but most likely to give the app access to their Bluetooth, location, and motion (see Fig. 9). This is not unexpected for a fitness tracker app.

In comparing against participants' existing settings, we note that Fitbit asks a different set of phone permissions depending on the user's mobile Operating System (OS). Among our respondents there are 162 iOS users, 103 Android 6.0+ users, 17 Android users with an older OS (which does not allow them to control each

[29] We ask for these preferences in our questionnaire because this information is mandatory in all fitness apps, hence, participants' actual disclosure does not necessarily reflect their true preferences.

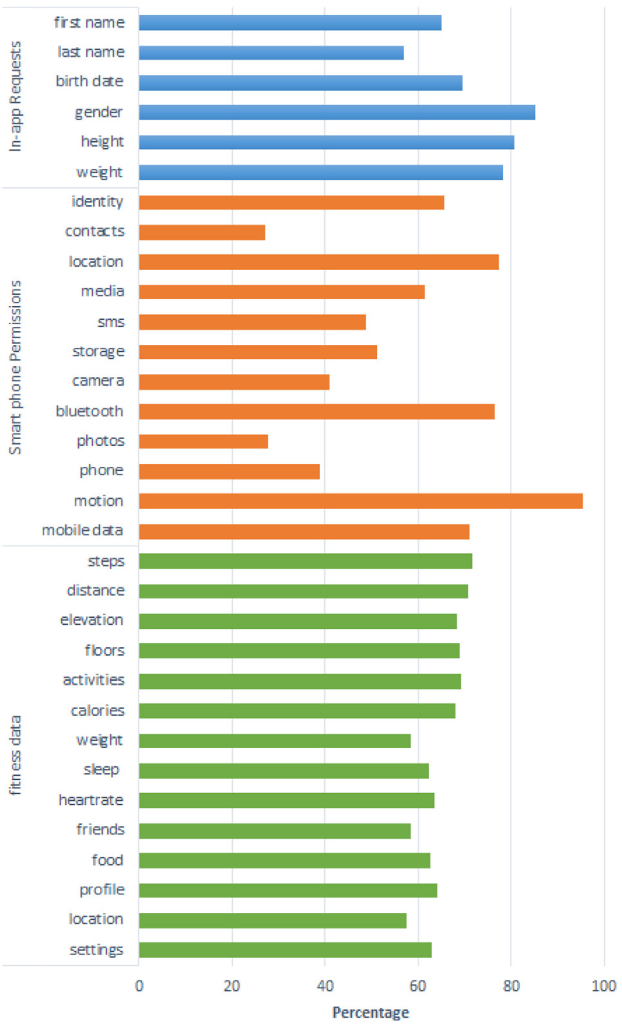

Fig. 9. FitPro permissions allowed by the participants.

permission separately), and 13 Windows users. In our evaluation we only consider the two larger groups of iOS and Android 6.0+ users.

Unlike the in-app requests, Android 6.0+ and iOS phone permissions are not mandatory, meaning users can allow or deny each permission separately. We asked participants to tell us their current permission settings for the Fitbit app. Given that participants can freely allow or deny each permission, we assume that their settings are aligned with their preferences. We compare these preferences with the PPM-based settings and the results show that they have significant statistical relationship for all permissions for both the Android 6.0+ and iOS, as shown in Table 2.

5.4.3. Fitness data

The fitness data produced by the user's fitness app can be accessed by other TPs to provide more services and features. There are, in fact, many of these external apps that use Fitbit's data. We ask participants to list the current permission settings of the external app they use most. Only 179 participants reported that they had an external app that accesses their fitness data, so for those who do not have an external app, we asked them what their preferences would be for sharing their fitness data with such an app. We report the results for these two groups of participants separately.

Note that in FitPro, exercise data are broken down into smaller granularity (i.e., steps, distance, floors, elevation, activity minutes, calories burned), giving users more options to control their privacy. However, we generalized these items into a single permission (i.e, Exercise) to be comparable with the settings on the external app that accesses their Fitbit's data.

For participants who have external apps, we compared the current settings of their most used third-party app with the PPM settings as shown in the left column of Table 3. It shows that their current permission settings show no association with their PPM settings (i.e., no statistical significance, all $p > 0.05$). It is possible, though, that their current settings do not reflect their real preferences. Indeed, the mismatch between users' privacy preferences and their settings is a phenomenon known as the ‘‘privacy paradox’’, which is well-established in previous research [70–74].

For participants who do not have third parties, we compared their preferences with their PPM settings, as shown in the right column of Table 3. It shows that their preferences are all significantly associated with their PPM settings. This means that participants' preferences on third-party sharing are captured by the PPM for all fitness data items ($p < 0.05$).

Another important thing shown in this figure is that the fitness data show very little variability (see Fig. 9). This is likely because *what* fitness data is shared is less important to the user than *who* the information is shared with (which is part of the GDPR permissions, as discussed below). This result is in accordance with previous studies such as in [75].

5.5. GDPR permissions

Unlike the permission sets discussed above, GDPR permissions are a novel contribution of our work, and hence we do not have participants' existing permissions to compare with. Hence, we simply report the results of the GDPR permission settings from our study in Fig. 10.

For the frequency of access, we let participants choose between granting FitPro continuous access, separate access for each workout (semi-continuous), or only grant access when using the app. Most participants only want to give access when using the app, which is a very useful privacy control since apps usually run on background even when not being used [76]. Those who chose to give the app continuous access may want their fitness tracker to count the calories burned and number of steps taken throughout the day, which is one of the main features of many fitness trackers.

For the retention of data, participants are given the following options: store until no longer used, store until the app is uninstalled, or store indefinitely (as to recover during app reinstallation). Only 1% of the participants chose the latter. Most of the participants prefer to retain their data until the app is uninstalled (47%) or would like to store it until no longer used (42%). This even split shows that participants have different preferences regarding retention, which means that this permission is important and must be controllable by the user.

The purposes of data collection are then specified. Among our participants, 85% allow data collection for health purposes, 84% for safety purposes, 54% for social purposes, 62% for convenience purposes, and only 17% for commercial purposes. Having the option to deny data use for commercial purposes could solve many privacy issues that stem from commercial disclosure without the user's informed consent. On the other hand, we acknowledge that this is an integral part of many companies' business model.

Finally, GDPR entity types include social media apps, fitness apps, commercial and government fitness programs, and other apps on the user's phone. Fitness tracker users mostly allow sharing to fitness apps and social apps, but the latter only if

Table 2
Chi-square tests of association between PPM settings and participants' preference on smartphone permissions.

Phone permissions	Current settings vs PPM settings	
(S set)	Android (6+)	iOS
Phone	8.2 ($p < 0.05$)	–
Storage	9.0 ($p < 0.05$)	–
SMS	17.9 ($p < 0.05$)	–
Contacts	14.6 ($p < 0.05$)	43.2 ($p < 0.05$)
Location	13.5 ($p < 0.05$)	33.8 ($p < 0.05$)
Camera	22.7 ($p < 0.05$)	17.6 ($p < 0.05$)
Bluetooth	–	26.0 ($p < 0.05$)
Photos	–	17.6 ($p < 0.05$)
Media & Music	–	33.6 ($p < 0.05$)
Motion & Fitness	–	10.8 ($p < 0.05$)
Mobile Data	–	37.2 ($p < 0.05$)

Table 3
The table of chi-squared test of association between PPM settings and participants' preferences on user fitness data.

Fitness data (F set)	Users w/ TPs: Current settings vs PPM settings	Users w/o TPs: Preference vs PPM settings
Exercise	0.1 ($p > 0.05$)	7.8 ($p < 0.05$)
Weight	0.2 ($p > 0.05$)	5.9 ($p < 0.05$)
Sleep	0.0 ($p > 0.05$)	6.9 ($p < 0.05$)
Heartrate	0.3 ($p > 0.05$)	15.0 ($p < 0.05$)
Food & Water	0.3 ($p > 0.05$)	12.5 ($p < 0.05$)
Location	0.7 ($p > 0.05$)	11.5 ($p < 0.05$)
Devices & Settings	0.5 ($p > 0.05$)	26.7 ($p < 0.05$)
Friends	0.6 ($p > 0.05$)	27.4 ($p < 0.05$)
Profile	1.0 ($p > 0.05$)	13.3 ($p < 0.05$)

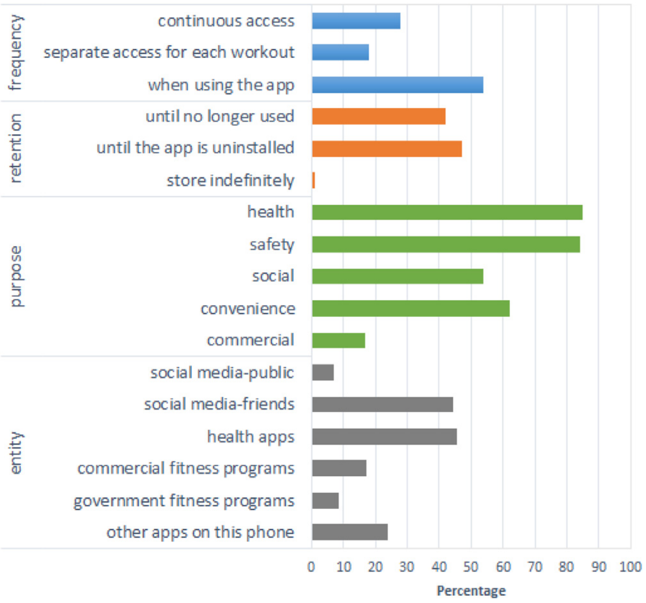

Fig. 10. GDPR permissions allowed by the users.

sharing can be restricted to their friends. In fact, participants are least likely to share their data on social media publicly.

In general, our study is among the first to measure users' preferences regarding GDPR-mandated permissions. Our results show a substantial variability in users' preferences regarding these permissions, which is a testament to the importance of these permissions in the fitness domain, and likely beyond as well.

6. Limitations and future work

The presented approach contributes to the research line that aims to support users in managing their privacy preferences in light of the continuing proliferation of IoT devices and TP apps. Differently from other approaches (see [77] for an overview on Privacy Engineering in the IoT), our proposal does not require that the user data are stored in the PDM—only the privacy preferences have to be defined and stored. We did not focus specifically on the means of user preference acquisition and storage—preferences can be requested using the PPM, or they can be imported or even inferred. We plan to address this aspect in future papers and we are currently working on this issue.

We also did not design PPIoT with a temporal context in mind (e.g., modeling changes in the user's preferences over time). This is also true for the existing ontology on which PPIoT is based (PPO). Following the approach used by the designers of the PPO ontology, we regard PPIoT as a vocabulary that can be used to express preferences, while privacy managers can handle the changes in privacy preferences and the implications of such changes. However, future work can also develop a temporal context for the PPIoT ontology itself.

An obvious further weakness of our approach lies in the lack of control over the data once it has been disclosed to the TP. Third-party access can be allowed or restricted via the condition properties, but there are no guarantees that these properties are respected by the TP. Extensive data sharing among TPs can also result in additional inference risks. As such, our framework only applies to situations where TPs can be trusted and/or held accountable for their actions—but this is also true for traditional TP policy statements.

Our proposal aims to automate the matching and negotiation between the TP policy statement and the user's preferences, resulting in a more transparent and controllable management of privacy permissions in the IoT context. This requires TPs to comply with the PPIoT Ontology. Note, though, that our approach uses "graceful degradation", where traditional privacy policies are automatically mapped to the PPIoT Ontology. The PDM can still operate in this case; only negotiation is not possible on the TPs part.

Our work can help the TPs abide by the GDPR regulations by making their policies more transparent and controllable, and by allowing the TP to acquire explicit consent from the user as discussed in Section 4.1. However, we concur that the PPM will likely not completely replace traditional policy statements, which are written in a certain way to provide legal protection. That said, we argue that our interactive privacy negotiation is a big step towards the newly implemented GDPR requirements.

Moreover, although the PPIoT Ontology is flexible and expressive, its effectiveness relies upon the proper identification and representation of preferences by the PDM—a difficult and often error-ridden task, especially for end users [78]. In future work, we therefore aim to automate the creation of user profiles.

Automating the PDM process is specifically important in light of its more ubiquitous implementation. In our current scenario of a single fitness tracking application, the PDM process happens in a single session. However, if the PDM concept is implemented more pervasively, then one can expect the PDM to provide users with more or less continuous privacy setting recommendations. The degree of automation and the dynamics of this process are important topics for future work, while the focus of this paper is on the interaction and data model that support the negotiation and the recommendation of privacy settings. The PPM workflow always ends with a recommendation, and thus a "manual" action from the user, in accordance with GDPR principles that require the user to give explicit consent to each request, but the PDM simplifies this task and, importantly, in the case of disagreement between the user privacy preference and TP request, it tries first to automatically negotiate with the enhanced TP, in order to favor the user's privacy preferences.

Finally, the PDM can only support negotiation when the TP's policy statement is encoded in PPIoT. When automatic negotiation is not possible, the TP may still allow for some manual settings to be made by the user. Arguably, users may find it difficult to set these settings on their own. Our future work will address how the PDM can give recommendations in such situations that are phrased in a way that optimizes the user's confidence (cf. [79]).

We are aware that ontology development is necessarily an iterative process, and this process of iterative design will likely continue through the entire life-cycle of the ontology, in accordance with [51]. Thus, it is necessary to have mechanisms to adapt the ontology to eventual domain changes and accordingly to facilitate the validation of these adaptations. This can be done through alignments to new or updated ontologies that might be developed, as discussed in the Related Work Section with regard to the new GDPR-based ontologies that are being published. A similar future realignment may be necessary for the IoT domain as well.

7. Contribution and conclusion

This paper presents an SWT-based solution for supporting IoT users in managing their privacy preferences. Our PPIoT ontology enables the representation of fine-grained privacy preferences and handles the complexity of heterogeneous personal IoT devices, while our PDM mediates the interaction between the user and TPs, enabling the negotiation of fine-grained privacy settings. Our main contribution to the literature is the combination and extension of previous approaches for SWT-based privacy management to cover the demands of the IoT domain, which is compliant with the FIP principles and the GDPR. Our framework also allows for negotiation on both the TP side and the user side, thereby balancing the privacy and utility of the service.

Below, we summarize some of the improvements we made compared to existing work.

- The PPIoT Ontology improves upon PPO by providing a richer preference model that fits in the IoT paradigm, thereby allowing users to create fine-grained privacy preferences [10].
- The ontology in [19] allows users to place regulations and conditions on factors based on the purpose of data recipient, usage and retention, disputes, remedy, and access control. This was included in the PPIoT extensions.
- Likewise, the PROACT Ontology [18] focuses on ubiquitous computing and includes a larger set of privacy concepts formalized in the notion of an "activity sphere". However, the abstract nature of this ontology makes it difficult to capture the complexity of the IoT paradigm. Arguably, our work captures a similar level of granularity.
- The ontology for privacy rules described in [25] addresses the privacy challenges of context-aware systems by defining a separate "data" class and a "condition" class. Our work adopts this approach, but also allows for *negotiation* of the conditions of each type of data between the user and the TP. We believe that this negotiation, in the context of an extensible SW Ontology, is crucial for the feasibility of privacy management in the context of IoT.

Finally, our PPIoT ontology and the interactive PPM help TPs meet the GDPR requirement of providing straightforward policy statements and requesting explicit consent.

Declaration of competing interest

No author associated with this paper has disclosed any potential or pertinent conflicts which may be perceived to have impending conflict with this work. For full disclosure statements refer to https://doi.org/10.1016/j.future.2019.10.024.

Acknowledgment

The research project was partially supported by National Science Foundation, USA award no. 1640664 received by Dr. Knijnenburg and a gift from Samsung Research America.

References

[1] É. Morin, M. Maman, R. Guizzetti, A. Duda, Comparison of the device lifetime in wireless networks for the internet of things, IEEE Access 5 (2017) 7097–7114.

[2] N. Heuveldop, et al., Ericsson mobility report, Tech. Rep. EAB-17 5964, Ericsson AB, Technol. Emerg. Business, Stockholm, Sweden, 2017.

[3] Statista, Forecast on Connected Devices Per Person Worldwide 2003–2020, 2018.

[4] M. Smith, C. Szongott, B. Henne, G. Von Voigt, Big Data privacy issues in public social media, in: 2012 6th IEEE International Conference on Digital Ecosystems and Technologies, DEST, IEEE, 2012, pp. 1–6.

[5] F. Carmagnola, F. Osborne, I. Torre, Escaping the big brother: An empirical study on factors influencing identification and information leakage on the Web, J. Inf. Sci. 40 (2) (2014) 180–197.

[6] T. Erl, W. Khattak, P. Buhler, Big Data Fundamentals: Concepts, Drivers & Techniques, Prentice Hall Press, 2016.

[7] P. Bahirat, Y. He, A. Menon, B. Knijnenburg, A data-driven approach to developing IoT privacy-setting interfaces, in: 23rd International Conference on Intelligent User Interfaces, IUI '18, ACM, Tokyo, Japan, 2018, pp. 165–176.

[8] M. Ganzha, M. Paprzycki, W. Pawłowski, P. Szmeja, K. Wasielewska, Semantic interoperability in the Internet of Things: an overview from the INTER-IoT perspective, J. Netw. Comput. Appl. 81 (2017) 111–124.

[9] M. Bermudez-Edo, T. Elsaleh, P. Barnaghi, K. Taylor, IoT-Lite: a lightweight semantic model for the internet of things and its use with dynamic semantics, Pers. Ubiquitous Comput. (2017) 1–13.

[10] O. Sacco, J.G. Breslin, PPO & PPM 2.0: Extending the privacy preference framework to provide finer-grained access control for the web of data, in: Proceedings of the 8th International Conference on Semantic Systems, 2012, pp. 80–87.

[11] F. Gilbert, Privacy and security legal issues, in: Internet of Things and Data Analytics Handbook, Wiley Online Library, 2017, pp. 699–718.

[12] The European Parliament and the Council of the European Union, Regulation (EU) 2016/679 Of The European Parliament and of The Council, Off. J. Eur. Union (2016) 1:88.

[13] A. Sunyaev, T. Dehling, P.L. Taylor, K.D. Mandl, Availability and quality of mobile health app privacy policies, J. Am. Med. Inform. Assoc. 22 (e1) (2014) e28–e33.

[14] European Commission, A new era for data protection in the EU; What changes after May 2018, 2016, pp. 1–3, URL https://ec.europa.eu/commission/sites/beta-political/files/data-protection-factsheet-changes_en.pdf.

[15] P. Beatty, I. Reay, S. Dick, J. Miller, P3p adoption on e-commerce web sites: A survey and analysis, IEEE Internet Comput. 11 (2) (2007) 65–71.

[16] H. Lee, A. Kobsa, Privacy preference modeling and prediction in a simulated campuswide iot environment, in: IEEE International Conference on Pervasive Computing and Communications, PerCom, IEEE, 2017, pp. 276–285.

[17] Y.-J. Hu, J.-J. Yang, A semantic privacy-preserving model for data sharing and integration, in: Proceedings of the International Conference on Web Intelligence, Mining and Semantics, ACM, 2011, p. 9.

[18] I. Panagiotopoulos, L. Seremeti, A. Kameas, V. Zorkadis, Proact: An ontology-based model of privacy policies in ambient intelligence environments, in: 14th Panhellenic Conference on Informatics, PCI, IEEE, 2010, pp. 124–129.

[19] P. Bodorik, D. Jutla, M.X. Wang, Consistent privacy preferences (cpp): model, semantics, and properties, in: Proceedings of the 2008 ACM Symposium on Applied Computing, ACM, 2008, pp. 2368–2375.

[20] J. Lopez, R. Rios, F. Bao, G. Wang, Evolving privacy: From sensors to the Internet of Things, Future Gener. Comput. Syst. 75 (2017) 46–57.

[21] A.M. Turri, R.J. Smith, S.W. Kopp, Privacy and RFID technology: A review of regulatory efforts, J. Consum. Aff. (2017).

[22] M. Compton, P. Barnaghi, L. Bermudez, R. GarcíA-Castro, O. Corcho, S. Cox, J. Graybeal, M. Hauswirth, C. Henson, A. Herzog, et al., The SSN ontology of the W3C semantic sensor network incubator group, Web Semant. Sci. Serv. Agents World Wide Web 17 (2012) 25–32.

[23] R. Atkinson, R. García-Castro, J. Lieberman, C. Stadler, Semantic Sensor Network Ontology W3C Recommendation, 2017.

[24] C. Perera, C. Liu, R. Ranjan, L. Wang, A.Y. Zomaya, Privacy-knowledge modeling for the Internet of Things: A look back, Computer 49 (12) (2016) 60–68.

[25] N. Zhang, C. Todd, Developing a Privacy Ontology for Privacy Control in Context-Aware Systems, Dept. of Electronic & Electrical Eng., Univ. College London, 2006.

[26] L. Kagal, T. Finin, M. Paolucci, N. Srinivasan, K. Sycara, G. Denker, Authorization and privacy for semantic web services, IEEE Intell. Syst. 19 (4) (2004) 50–56.

[27] Z. Iqbal, J. Noll, S. Alam, M.M. Chowdhury, Toward user-centric privacy-aware user profile ontology for future services, in: Third International Conference on Communication Theory, Reliability, and Quality of Service, 2010, pp. 249–254.

[28] L.A.F. Martimiano, M.R.P. Goncalves, E. dos Santos Moreira, An ontology for privacy policy management in ubiquitous environments, in: Network Operations and Management Symposium, NOMS 2008, IEEE, 2008, pp. 947–950.

[29] H. Benfenatki, F. Biennier, M. Winckler, L. Goncalves, O. Nicolas, Z. Saoud, Towards a User Centric Personal Data Protection Framework, in: ACM CHI Conference on Human Factors in Computing Systems, 2018.

[30] S. Foresti, P. Samarati, Supporting user privacy preferences in digital interactions, in: Computer and Information Security Handbook, Elsevier, 2017, pp. 801–822.

[31] M. Bennicke, et al., Towards automatic negotiation of privacy contracts for internet services, in: The 11th IEEE International Conference on Networks, 2003. ICON2003, IEEE, 2003, pp. 319–324.

[32] H. Li, D. Ahn, P.C. Hung, Algorithms for automated negotiations and their applications in information privacy, in: Proceedings. IEEE International Conference on E-Commerce Technology, 2004. CEC 2004, IEEE, 2004, pp. 255–262.

[33] I. Jang, H.S. Yoo, Personal information classification for privacy negotiation, in: 2009 Fourth International Conference on Computer Sciences and Convergence Information Technology, IEEE, 2009, pp. 1117–1122.

[34] I.J. Jang, W. Shi, H.S. Yoo, Policy negotiation system architecture for privacy protection, in: 2008 Fourth International Conference on Networked Computing and Advanced Information Management, Vol. 2, IEEE, 2008, pp. 592–597.

[35] S.-C. Cha, M.-S. Chuang, K.-H. Yeh, Z.-J. Huang, C. Su, A user-friendly privacy framework for users to achieve consents with nearby BLE devices, IEEE Access 6 (2018) 20779–20787.

[36] S.-C. Cha, T.-Y. Tsai, W.-C. Peng, T.-C. Huang, T.-Y. Hsu, Privacy-aware and blockchain connected gateways for users to access legacy IoT devices, in: 2017 IEEE 6th Global Conference on Consumer Electronics, GCCE, IEEE, 2017, pp. 1–3.

[37] R. Aydoğan, P. Øzturk, Y. Razeghi, Negotiation for incentive driven privacy-preserving information sharing, in: International Conference on Principles and Practice of Multi-Agent Systems, Springer, 2017, pp. 486–494.

[38] K. Alanezi, S. Mishra, A privacy negotiation mechanism for the internet of things, in: 2018 IEEE 16th Intl Conf on Dependable, Autonomic and Secure Computing, 16th Intl Conf on Pervasive Intelligence and Computing, 4th Intl Conf on Big Data Intelligence and Computing and Cyber Science and Technology Congress, DASC/PiCom/DataCom/CyberSciTech, IEEE, 2018, pp. 512–519.

[39] T. Baarslag, A.T. Alan, R. Gomer, M. Alam, C. Perera, E.H. Gerding, et al., An automated negotiation agent for permission management, in: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, International Foundation for Autonomous Agents and Multiagent Systems, 2017, pp. 380–390.

[40] R. Gellman, Fair information practices: A basic history, SSRN (2017).

[41] S. Landau, What was samsung thinking? IEEE Secur. Priv. 13 (3) (2015) 3–4.

[42] H.J. Pandit, K. Fatema, D. O'Sullivan, D. Lewis, Gdprtext-GDPR as a linked data resource, in: European Semantic Web Conference, Springer, 2018, pp. 481–495.

[43] ICO UK, Getting ready for the GDPR, 2018, https://ico.org.uk/for-organisations/data-protection-self-assessment/, (Online; accessed 10-April-2019).

[44] Microsoft Trust Center, Detailed GDPR assessment, 2018, http://aka.ms/gdprdetailedassessment/, (Online; accessed 10-April-2019).

[45] S. Agarwal, S. Steyskal, F. Antunovic, S. Kirrane, Legislative compliance assessment: Framework, model and GDPR instantiation, in: Annual Privacy Forum, Springer, 2018, pp. 131–149.

[46] M. Palmirani, M. Martoni, A. Rossi, C. Bartolini, L. Robaldo, PrOnto: Privacy ontology for legal reasoning, in: International Conference on Electronic Government and the Information Systems Perspective, Springer, 2018, pp. 139–152.

[47] M. Palmirani, G. Governatori, Modelling legal knowledge for GDPR compliance checking, Frontiers Artificial Intelligence Appl. 313 (2018) 101–110.

[48] L. Elluri, K.P. Joshi, et al., A knowledge representation of cloud data controls for EU GDPR compliance, in: 11th IEEE International Conference on Cloud Computing, CLOUD, 2018.

[49] H. Pandit, D. Lewis, Modelling provenance for GDPR compliance using linked open data vocabularies, in: 5th Workshop on Society, Privacy and the Semantic Web - Policy and Technology, PrivOn 2017, CEUR 1951 (2017), 2017, URL http://ceur-ws.org/Vol-1951/#paper-06.

[50] I. Torre, O.R. Sanchez, F. Koceva, G. Adorni, Supporting users to take informed decisions on privacy settings of personal devices, Pers. Ubiquitous Comput. 22(2) (2018) 345—364.

[51] N.F. Noy, D.L. McGuinness, et al., Ontology development 101: A guide to creating your first ontology, technical report KSL-01-05 and ..., Stanford knowledge systems laboratory, 2001.

[52] R. Studer, V. Benjamins, D. Fensel, Knowledge engineering: principles and methods, Data Knowl. Eng. 25 (1–2) (1998) 161–197.

[53] M. Noura, A. Gyrard, S. Heil, M. Gaedke, Concept extraction from the Web of Things knowledge bases, in: Proceedings of the International Conference WWW/Internet, 2018.

[54] M.A. Musen, et al., The protégé project: a look back and a look forward, AI Matters 1 (4) (2015) 4.

[55] H. Hlomani, D. Stacey, Approaches, methods, metrics, measures, and subjectivity in ontology evaluation: A survey, Semant. Web J. 1 (5) (2014) 1–11.

[56] A. Gómez-Pérez, Ontology evaluation, in: Handbook on Ontologies, Springer, 2004, pp. 251–273.

[57] D. Vrandečić, Ontology evaluation, in: Handbook on Ontologies, Springer, 2009, pp. 293–313.

[58] D. Tsarkov, I. Horrocks, FaCT++ description logic reasoner: System description, in: International Joint Conference on Automated Reasoning, Springer, 2006, pp. 292–297.

[59] R. Shearer, B. Motik, I. Horrocks, HermiT: A highly-efficient OWL reasoner, in: OWLED, Vol. 432, 2008, p. 91.

[60] E. Sirin, B. Parsia, Pellet: An owl dl reasoner, in: Proc. of the 2004 Description Logic Workshop, DL 2004, 2004, pp. 212–213.

[61] J. Brank, M. Grobelnik, D. Mladenić, A Survey of Ontology Evaluation Techniques, 2005.

[62] O. Sacco, A. Passant, A privacy preference manager for the social semantic Web, in: SPIM, 2011, pp. 42–53.

[63] R. Porzel, R. Malaka, A task-based approach for ontology evaluation, in: ECAI Workshop on Ontology Learning and Population, Valencia, Spain, Citeseer, 2004, pp. 1–6.

[64] T. Shah, F. Rabhi, P. Ray, K. Taylor, enhancing automated decision support across medical and oral health domains with semantic web technologies, in: Proceedings of the 24th Australasian Conference on Information Systems, 2014.

[65] O.R. Sanchez, I. Torre, H. Yangyang, B. Knijnenburg, A recommendation approach for user privacy preferences in the fitness domain, User Model. User-Adapt. Interact. (2019) http://dx.doi.org/10.1007/s11257-019-09246-3.

Please cite this article as: O.R. Sanchez, I. Torre and B.P. Knijnenburg, Semantic-based privacy settings negotiation and management, Future Generation Computer Systems (2019), https://doi.org/10.1016/j.future.2019.10.024.

[66] I. Torre, F. Koceva, O.R. Sanchez, G. Adorni, A framework for personal data protection in the IoT, in: 2016 11th International Conference for Internet Technology and Secured Transactions, ICITST, IEEE, 2016, pp. 384–391.

[67] R. Yus, E. Mena, Mobile endpoints: Accessing dynamic information from mobile devices, in: International Semantic Web Conference (Posters & Demos), 2015.

[68] L. Tsai, P. Wijesekera, J. Reardon, I. Reyes, S. Egelman, D. Wagner, N. Good, J.-W. Chen, Turtle guard: Helping Android users apply contextual privacy preferences, in: Symposium on Usable Privacy and Security, SOUPS, 2017.

[69] C. Jensen, C. Potts, Privacy policies as decision-making tools: an evaluation of online privacy notices, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM, 2004, pp. 471–478.

[70] P.A. Norberg, D.R. Horne, D.A. Horne, The privacy paradox: personal information disclosure intentions versus behaviors, J. Consum. Aff. 41 (1) (2007) 100–126, http://dx.doi.org/10.1111/j.1745-6606.2006.00070.x, URL http://onlinelibrary.wiley.com/doi/10.1111/j.1745-6606.2006.00070.x/abstract.

[71] T. Dinev, P. Hart, An extended privacy calculus model for e-commerce transactions, Inf. Syst. Res. 17 (1) (2006) 61–80, http://dx.doi.org/10.1287/isre.1060.0080, URL http://isr.journal.informs.org/cgi/content/abstract/17/1/61.

[72] C. Van Slyke, J.T. Shim, R. Johnson, J.J. Jiang, Concern for information privacy and online consumer purchasing, J. Assoc. Inf. Syst. 7 (1) (2006) URL http://aisel.aisnet.org/jais/vol7/iss1/16.

[73] S. Spiekermann, J. Grossklags, B. Berendt, Stated privacy preferences versus actual behaviour in EC environments: A reality check, in: WI-if 2001: The 5th International Conference Wirtschaftsinformatik - 3rd Conference Information Systems in Finance, Augsburg, Germany, 2001, pp. 129–148.

[74] H. Xu, X.R. Luo, J.M. Carroll, M.B. Rosson, The personalization privacy paradox: An exploratory study of decision making process for location-aware marketing, Decis. Support Syst. 51 (1) (2011) 42–52, http://dx.doi.org/10.1016/j.dss.2010.11.017, ACM ID: 1943793.

[75] H. Lee, A. Kobsa, Privacy preference modeling and prediction in a simulated campuswide iot environment, in: 2017 IEEE International Conference on Pervasive Computing and Communications, PerCom, IEEE, 2017, pp. 276–285.

[76] P. Wijesekera, A. Baokar, A. Hosseini, S. Egelman, D. Wagner, K. Beznosov, Android permissions remystified: A field study on contextual integrity, in: 24th USENIX Security Symposium, USENIX Security 15, 2015, pp. 499–514.

[77] A. Kung, F. Kargl, S. Suppan, J. Cuellar, H.C. Pöhls, A. Kapovits, N.N. McDonnell, Y.S. Martin, A privacy engineering framework for the Internet of Things, in: Data Protection and Privacy: (In) Visibilities and Infrastructures, Springer, 2017, pp. 163–202.

[78] B.P. Knijnenburg, A. Kobsa, Making decisions about privacy: information disclosure in context-aware recommender systems, ACM Trans. Interact. Intell. Syst. 3 (3) (2013) 20:1–20:23.

[79] M. Namara, H. Sloan, P. Jaiswal, B.P. Knijnenburg, The potential for user-tailored privacy on Facebook, in: IEEE Symposium on Privacy-Aware Computing, Washington, D.C., 2018.

Odnan Ref D. Sanchez received BS.c. degree from the CIT University, Philippines in 2011 and M.Sc. degree from the University of Genoa, Italy in 2015. He is currently a Ph.D. student in Digital Humanities at the University of Genoa, Italy. His research interests include Machine learning, Privacy, and IoT systems.

Ilaria Torre, Ph.D. is a Professor at the Department of Informatics, Bioengineering, Robotics and Systems Engineering, University of Genova and a member of the Advisory Board of the Ph.D. School in Digital Humanities. Her main research interests include intelligent user interfaces, recommender systems, ubiquitous interactive systems, IoT and Semantic Web. She published over 70 papers in international journals and conferences. In recent years, she organized and co-organized several international events, including the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP 2020), the ACM Intelligent User Interfaces Student Consortium (IUI SC 2019), and the Internet of Things for Active and Assisted Living workshop at the IEEE GIobal IoT Summit 2017 (IoTAAL 2017).

Bart Knijnenburg is an Assistant Professor in Human-Centered Computing at the Clemson University School of Computing where he co-directs the Humans and Technology lab. He holds a B.S. in Innovation Sciences and an M.S. in Human-Technology Interaction from the Eindhoven University of Technology, The Netherlands, an M.A. in Human–Computer Interaction from Carnegie Mellon University, and a Ph.D. in Information and Computer Sciences from UC Irvine. Bart works on privacy decision-making and user-centric evaluation of adaptive systems. His research has received funding from the National Science Foundation, the Department of Defense, and corporate sponsors.

Please cite this article as: O.R. Sanchez, I. Torre and B.P. Knijnenburg, Semantic-based privacy settings negotiation and management, Future Generation Computer Systems (2019), https://doi.org/10.1016/j.future.2019.10.024.