

Personal Data Privacy Semantics in Multi-Agent Systems Interactions

Davide Calvaresi¹, Michael Schumacher¹, and Jean-Paul Calbimonte¹

University of Applied Sciences and Arts Western Switzerland (HES-SO),
Sierre, Switzerland
`{firstname.lastname}@hevs.ch`

Abstract. In recent years, we have witnessed the growth of applications relying on the use and processing of personal data, especially in the health and well-being domains. Users themselves produce these data (e.g., through self-reported data acquisition, or personal devices such as smartphones, smartwatches or other wearables). A key challenge in this context is to guarantee the protection of personal data privacy, respecting the rights of users for deciding about data reuse, consent to data processing and storage, anonymity conditions, or the right to withhold or delete personal data. With the enforcement of recent regulations in this domain, such as the GDPR, applications are required to guarantee compliance, challenging current practices for personal data management. In this paper, we address this problem in the context of decentralized personal data applications, which may need to interact and negotiate conditions of data processing and reuse. Following a distributed paradigm without a top-down organization, we propose an agent-based model in which personal data providers and data consumers are embedded into privacy-aware agents capable of negotiating and coordinating data reuse, consent, and policies, using semantic vocabularies for privacy and provenance.

Keywords: Privacy ontologies · Agent data privacy · Semantic agents.

1 Introduction

Protecting data privacy and complying with privacy policies is of utmost importance, especially when handling sensitive personal information. Beyond personal datasets, including demographic information, or medical history records, nowadays, the digitization era has opened the way for a large number of data acquisition alternatives. Ranging from applications installed in smartphones to well-being sensors embedded in smartwatches, or social network data collected on our behalf, the amount of sensitive information directly or indirectly collected by users and third-parties has experienced enormous growth. This trend entails several ethical and legal challenges, which cannot be isolated from the technological implications behind the acquisition of these datasets [3].

The European Union introduced and adopted in 2018 the General Data Protection Regulation (GDPR), a comprehensive legislation body that has had

an enormous impact on how personal data is collected, stored, processed, and shared [26]. The legal enforcement of GDPR is also transforming how digital solutions, applications, and systems handle sensitive data. For instance, the need for explicit consent for the use of data, the right to timely receive all data collected for oneself, or the right to completely delete personal data are specified in this law, forcing technical solutions to be provided following the regulations. Although many privacy-preserving and GDPR compliant frameworks have been designed and implemented in the last few years [2], most of these rely on centralized enforcement and compliance, often assuming full control inside the boundaries of a silo-ed system. For instance, this is the case with clinical studies, in which acquisition instruments, such as wearable devices, have built-in restrictions to guarantee compliance. This is also the case within the boundaries of a hospital, in which both monitoring devices and electronic health records are designed to respect privacy regulations. However, these approaches disregard the decentralized nature of data interactions in larger scope scenarios. Having full control in a top-down fashion is not feasible in more complex environments such as crowd-sourcing, or public data collection in which no central authority can be the sole entity in charge of data protection. Furthermore, even when centralized data privacy is enforced, the representation of data privacy needs, consent, purpose, etc. lacks the expressiveness and machine-understandable semantics required to provide automated management of personal data handling.

In this paper, we propose the adoption of decentralized agent-based data privacy negotiation, coordination, and enforcement, using semantic representations of personal data privacy conditions and handling. More precisely, we define a set of minimal personal data privacy interaction requirements among agents and the design principles of privacy-aware agent interactions regarding personal data handling. Then, we propose a conceptual architecture in which these interactions are translated into multi-agent protocol specifications, annotated with semantic information about the purpose, recipient, processing, and consent of the personal data. We base this specification in the Data Privacy Vocabulary (DPV) [19], developed by the Data Privacy Vocabularies and Controls Community Group, hosted under the umbrella of the W3C organization.

The remainder of the paper is organized as follows: Section 2 provides a motivating use-case and elaborates requirements. Section 3 presents the design principles of a MAS personal data privacy. Section 4 elaborates on data privacy semantics in agent(s) interactions. Section 5 summarizes the related work. Finally, Section 6 concludes the paper and proposes a road-map for future steps.

2 Personal Data Privacy Interactions

When sharing, reusing, or processing personal data, privacy concerns surface at different stages, and specific regulations need to be considered at each of them. For example, Figure 1 considers a use case in physical rehabilitation, where motion monitoring sensors are used to track exercise and physical activity from a patient [4, 5]. Following a traditional approach, the patient will sign a general

consent for data collection and processing, which will be managed by the clinic or hospital to which the physiotherapist is affiliated. After this consent is signed, the wearable sensors can collect data that will be forwarded to the monitoring system managed by the physiotherapist. In turn, such data can also be linked to the electronic health record of the patient, and later shared with other healthcare professionals within the hospital or clinic.




Fig. 1: Requests from a physiotherapist to patients’ data, according to a general consent. This implies no customization of consent conditions, and disregards the decentralized nature of sensing applications.

In this scenario, the patient (or subject) has granted permission to several actions and activities, many of which are probably not exactly clear or transparent. Therefore the following questions can be raised:

- Can the subject timely access all collected data during the interventions?
- Is the subject able to opt-out of specific processing/monitoring activities?
- Can the subject establish restrictions on types of data to be collected/reused?
- How can the subject trace the actions and data access of healthcare providers?
- Can the subject limit read/write access to specific healthcare providers?
- Is the subject given the possibility of deleting or withdrawing her data completely or partially?
- Can the user dynamically change the consent conditions, including restrictions on specific data handling purposes?
- Can the subject be notified of risks or evidence of privacy breach or other undesired activities?

Most of these questions can be linked to regulations in data protection laws. In particular, for the EU, GDPR precisely establishes a legal framework to guarantee that subjects can have satisfactory answers to all these questions. For instance, concerning the possibility of deleting one’s data, GDPR introduces the right to be forgotten [26]: *“The data subject shall have the right to obtain from the controller the erasure of personal data without undue delay”*. In the use-case mentioned above, the broad consent given by the patient leaves little space for fine-grained management of personal data, resulting in activities (e.g., processing, data reuse, profile learning) which will be somehow hidden to the patient. Although this is generally based on the assumption of a trusted relationship, nowadays people are more and more aware of the importance of privacy, and on the potential benefits of having access not only to one’s data but also to the

trace of usage of that data. Moreover, given the potential complexity of data privacy regulations, people may often require assistance to fully grasp the implications of individual consent grant decisions. At the same time, it is also needed to provide users with the means to keep track of data usage and risks, while allowing integration even among different healthcare institutions. Considering that in many cases, patients may desire to share personal data among different clinics and hospitals, their privacy preferences should be able to be transmitted and enforced across institutional and administrative boundaries.

Taking into account these considerations, we formulate the following requirements for establishing personal data privacy in a decentralized environment:

- R1: *Data handling actors*: Before establishing any privacy interactions, a shared understanding of actors and processes for data handling must be established. This model, following legal regulations such as the GDPR, must formulate who are the data controllers, subjects, recipients, and what are the possible data handling processes that they may activate. These actors must be able to establish their own goals with respect to the data (e.g., compliance, privacy policies), have their own knowledge or state, as well as a set of potential actions or intentions.
- R2: *Decentralized interactions*: The different actors in charge of personal data handling must specify the possible interactions among them, without the need of a centralized entity governing their decisions. Considering that a data controller (e.g., sensor data collection on behalf of a clinical provider) may need to request consent acceptance to a subject (e.g., a patient), none of these actors should be imposed decisions regarding the negotiation of the data access conditions nor should be able to take them independently.
- R3: *Semantic data privacy modelling*: To have meaningful interactions among data privacy controllers and subjects, it is essential to rely on the standard and human/machine-understandable models that represent data handling purposes, processes, consent, and other data privacy characteristics. These models must be specified using semantic representations, which can embed interpretable logic that can be later used for enforcing data privacy policies. Semantic vocabularies —aligned with current legislation such as the GDPR— are required to attain this degree of interpretability.
- R4: *Interaction protocols*: The interactions must follow a well-defined pattern, specified as a set of behaviors, so that they allow the negotiation or collaboration among different entities. For example, if a clinical study requires crowd-sourcing personal data from a given population, a surrogate entity may emit a call-for-data, including consent and policy conditions, to which potential data providers could emit “accept” or “reject” interactions, followed –if positive– by periodic data collection messages.
- R5: *Legal compliance*: All interactions among entities dealing with personal data must comply with the applicable legal framework, e.g., GDPR in the EU.

Furthermore, once the interaction model has been established under these conditions, it is also important to enforce the following aspects, directly regarding the handling of personal data privacy:

- R6: *Verification*: It should be possible that all entities participating in data privacy interactions can verify the compliance to regulations. This verification should be automatized, even if across institutional boundaries, thanks to the semantic representation of policies and handling conditions.
- R7: *Tracking*: It must be possible to keep track of all interactions, as well as reuse, access, processing, and handling events across the lifetime of personal datasets.
- R8: *Explainability*: Controllers should expose explainable and understandable interfaces for all data handling processes. This should allow users and subjects in general to have a clear understanding of data workflows and implications in privacy.
- R9: *Transparency*: Controllers should be able to timely communicate any event of importance to subjects, concerning data privacy, such as risks, breaches, compromises, or any other potentially relevant circumstance.
- R10: *Granularity*: It must be possible to choose the granularity at which personal data handling is performed. This includes the ability to select the purpose(s) for which data processing is requested, who has access and under which conditions, what are the potential data recipients, what type of reuse or publication is permitted, which technical measures will be applied, such as storage means, deadlines, etc.

3 Design Principles for MAS Personal Data Privacy

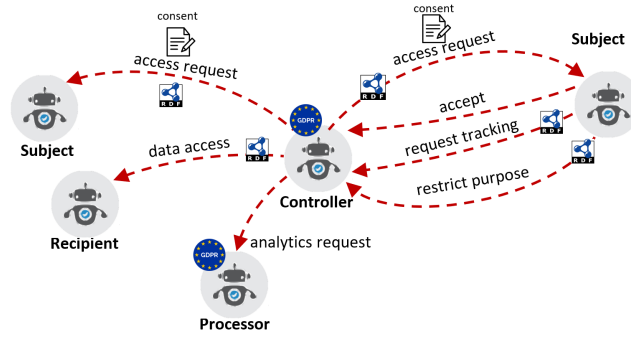


Fig. 2: Personal Data Privacy Agents. The controller may emit access request to subject agents. Negotiation, including consent may happen through agent interactions.

Considering the challenges and requirements enumerated previously, we propose a set of principles for establishing decentralized data privacy interactions among personal data providers, recipients, and managers. This set of principles is based on the adoption of the multi-agent system (MAS) paradigm, which has several properties that match the challenges of complex personal data exchange and reuse. First, the notion of MAS already implies the necessary degree of autonomy for agents, which can embody different types of data handling entities. Second, it naturally allows agents to set their own goals, which in terms of data

privacy can include specific policies, consent conditions, red lines regarding privacy, etc. Similarly, for data controllers, it may allow defining goals regarding the quality/quantity of data or de-anonymization guarantees. A third aspect refers to the possibility of establishing negotiating protocols and collaboration patterns among agents, which may translate to consent requests, data tracking petitions, right to be forgotten enforcement, etc. Finally, the agent paradigm permits the exchange of common knowledge, or beliefs, which can be crucial for personal data handling in complex scenarios, thanks to the usage of standard and semantically rich ontologies representing data privacy specifications. We identify three main design principles detailed in the following, and partially illustrated in Figure 2.

Decentralized agents: All participating entities in personal data privacy interactions are modelled as autonomous intelligent agents. Following the nomenclature of the GDPR we identify four main types of actors: *data controllers*, *subjects*, *recipients*, and *processors*. Controllers refer to people, organizations, or authorities that govern and decide about the purpose and processing of personal data. Subjects are the persons to which the data is related, while recipients are the people or entities to which the personal information is disclosed. Processors are those persons or entities that perform any processing of the personal data on behalf of the controller. Other classes of agents may exist, such as third parties or authorities, complementary to the main four (Figure 2). Each of these agents has its own set of goals, w.r.t. personal data handling. For example, a patient may require that all data that is shared with other agents should be only for academic research, or that it should be fully anonymized, or that it should exclude any profiling processing activities, etc. Notice that even under anonymization, re-identification is still possible through combination of different data sources, and agents may consider modelling potential attacks and contemplate countermeasures. These agents also have their own knowledge or beliefs, which may include metadata regarding the personal datasets under their control (e.g., for the data controllers), or the tracking activities of personal data (e.g., for a data subject). The agent knowledge can be arbitrarily complex, and it is the agent who decides which elements of it can be shared with other agents, and for which purposes.

Shared semantic vocabulary: The semantic interoperability among these agents is dictated by the use of a common ontology (or set of ontologies) establishing a common model for representing privacy data. This is a fundamental principle for the establishment of meaningful interactions among decentralized agents, given that there is not necessarily a sole authority governing the agent requests and responses. In this work, we advocate the use of the Data Privacy Vocabulary (DPV)¹, an ontology developed by the Data Privacy Vocabularies and Control Community Group, under the scope of the W3C (see Figure 3). This vocabulary, although not yet published as a standard, is a GDPR-based model supported by a group of academic, industrial, and administrative institutions, with a high

¹ <https://www.w3.org/ns/dpv>

potential for adoption in a wider scope [19]. The model vocabulary includes the definition of the main concepts regarding personal data handling, including consent, purpose, processing, legal basis, controllers, and recipients, among others.


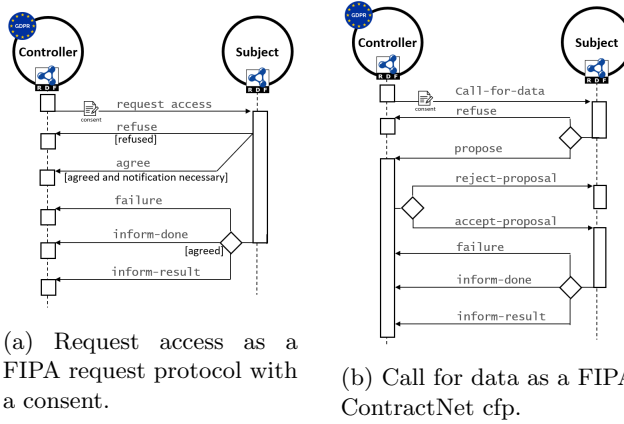


Fig. 3: Main classes of the Data Privacy Vocabulary [19]

Data privacy agent interactions: Having defined the agents and the semantics of the data that sets personal data privacy policies, the third main aspect refers to the specification of interactions in this context. In principle, we base the definition of these interactions in existing FIPA protocols. For instance, a data consent request can be embedded in a request interaction protocol, or a data crowd-sourcing request can be represented as a ContractNet protocol — thus, extending current approaches such as [15] which leverages on weighted aggregation of the encrypted users’ data via homomorphic cryptosystems or applications for mobile crowdsensing such as CarTel, ParkNet, BikeNet, and DietSense [8].



(a) Request access as a FIPA request protocol with a consent.

(b) Call for data as a FIPA ContractNet cfp.

Fig. 4: Agent data privacy interactions as FIPA protocols.

We identify a non-comprehensive minimum set of interactions listed below:

- Controller requests personal data (with consent) to a specific subject.
- Subject provides personal data (with a consent granted).
- Controller calls for personal data to a set of individuals represented by their subject agents.
- Subject selects only a certain purpose for data handling.
- Subject rejects request for data.
- Subject grants access to personal data only for a certain purpose.

- Subject/Controller customizes permissions and access rights.
- Controller tracks personal data reuse and processing.
- Subject deletes or withholds own personal data.
- Subject/controller verifies personal data use and policy
- Subject objects to data reuse or processing.
- Subject requests access to own personal data collected (and metadata)
- Controller notifies about data breaches or risk.

4 Data Privacy Semantics in Agent Interactions

As specified in the previous section, the semantic representation of personal data privacy information provides the foundations for interactions among heterogeneous agents. Regarding the specification of the interactions, these can be embedded into standard FIPA protocols².

Data controllers, generally in charge of data handling activities, have the possibility of requesting personal data from a subject, providing data to a recipient, initiating data collection, establishing and requesting consent proposals, requiring access rights, verifying policies, requesting processing activities, etc. Similarly, a data subject agent can request personal data tracking results, withhold personal data, reject access requests, choose which data purposes to apply to, request access to all collected data, etc. All of these interactions can be semantically represented, using models such as RDF, which allows the representation of information as triples (subject, predicate, object).

For example, in Listing 1.1, we illustrate how a FIPA-based interaction can be encoded in JSON-LD format. This interaction, based in the ContractNet protocol represents a request for data from a controller, to which different agents representing data subjects can answer through a bid. The message content, in this case, is a reference to a consent that subjects would need to agree with (in case of acceptance to participate in the data collection).

```
{ "prov:generatedAtTime": "2020-02-01T04:00:00.000Z",
  "@id": "ex:callForActivityData",
  "@graph": [
    { "@id": "ex:callForData1",
      "ag:permormative": "ag:CallForProposals",
      "ag:sender": "ex:controller1",
      "ag:protocol": "ag:ContractNet",
      "ag:ontology": "http://w3id.org/ns/dpv#",
      "ag:content": "ex:consentPatient1"  } ] }
```

Listing 1.1: Call for data representation in RDF JSON-LD format

All personal data handling activities, such as processing, data access request, consent management, etc., can also be represented using RDF. In Listing 1.2, we provide an example of data collection represented in RDF. Following the DPV ontology, it specifies the data controller (e.g., a hospital), the subject (e.g., a specific patient). It also indicates that the data to collect is about physical health; the purpose is for academic research and includes a consent.

² <http://www.fipa.org/repository/>


```

ex:dataRequest a dpv:PersonalDataHandling ;
  dpv:hasDataSubject      ex:patient1 ;
  dpv:hasPurpose          [a dpv:AcademicResearch] ;
  dpv:hasProcessing       [a dpv:Collect];   dpv:hasLegalBasis [a dpv:Consent];
  dpv:hasDataController  ex:hospital1;     dpv:hasRecipient    ex:physician3;
  dpv:hasPersonalDataCategory [a dpv:PhysicalHealth];
  dct:terms:title        "Personal Data Collection for clinical study ..." .

```

Listing 1.2: Data handling represented in RDF Turtle format.

Regarding the consent itself, it follows a similar structure as any data handling. In the example of Listing 1.3, the consent establishes three different purposes for data analysis: academic research, economic research, and personalized recommendations. Indeed, for instance, in a clinical study, analytics can be performed for research, but also to create recommendations that would benefit the patient. However, by examining this consent, the subject agent may choose only to authorize the analysis of data for academic research, banning the use for any commercial purpose. This enables fine-grained control over his/her own data.

```

ex:consentPatient1 a dpv:Consent ;
  dpv:hasDataSubject ex:patient1 ;
  dpv:hasPurpose     [a dpv:AcademicResearch], [a dpv:CommercialResearch],
                    [a dpv:CreatePersonalizedRecommendations] ;
  dpv:hasProcessing  [a dpv:Analyse];
  dct:terms:title   "Consent for Health data analysis in a clinical study ..." ;
  dpv:hasDataController ex:hospital1;
  dpv:hasRecipient    ex:physiotherapist1;
  dpv:hasPersonalDataCategory [a dpv:PhysicalHealth].

```

Listing 1.3: Consent represented in RDF Turtle format.

About the processing of data, the provenance ontology (PROV-O) provides a complementary set of classes and properties, which allows providing details about different types of data transformation activities. In Listing 1.4, we include an example of a data analytics task for personal data of a patient. The trace of the data analytics enables linking original datasets with processed results, which later can be requested by the data subject.

```

ex:dataAnalysis a dpv:Analysis ;
  dpv:hasDataSubject      ex:patient1 ;
  prov:used               ex:patientDataset1 ;
  dct:terms:title        "Data Analysis activity for patient data ..." ;
  prov:isAssociatedWith   ex:dataScientist1;
  prov:wasStartedAtTime  "2020-01-11T04:00:00.000Z".
ex:analyticsResults a prov:Entity ;
  prov:wasGeneratedBy    ex:dataAnalysis;   prov:wasDerivedFrom ex:patientDataset1.

```

Listing 1.4: Data analysis represented in RDF Turtle format.

5 Related Work

A large number of previous works have addressed the research problem of integrating privacy into multi-agent systems [16, 23]. Although multiple challenges have been discussed, studied, and addressed in these works, an analysis of the state of the art highlights that the inclusion of clear semantics for representing

personal data handling has constantly been missing. Even if several ontologies in this domain have been proposed [19], the inclusion of these within the context of applicable legislation (e.g., GDPR) and the interactions among agents has never been considered before. While some architectural propositions and even full implementations have incorporated agents as a central component of privacy-aware systems [20, 27, 10, 13], these generally lack the capability of having defined clear interactions in a decentralized manner, even across heterogeneous systems. Some other aspects have also been explored (i.e., the verification of privacy policies [14], or the establishment of transparent tracking of provenance [11, 25]). Many of these efforts are complementary to our approach, which could benefit from trust and explainability mechanisms [6, 7, 12], which have shown to be relevant to satisfy high standards in privacy regulations. Table 1 summarizes previous works wrt. the previously discussed aspects.

Aspect	Maturity Level		
	Conceptual	Implementation	Validation
<i>Agent architecture:</i> [20, 27, 16]	+++	+	+
<i>Decentralized interactions:</i> [10, 13]	+	+	-
<i>Semantic modelling:</i> [9, 21, 24, 19]	++	-	-
<i>Interaction protocols:</i> [1, 17, 22]	+	-	-
<i>Legal compliance:</i> [18, 28]	+	+	-
<i>Verification:</i> [14]	++	+	-
<i>Tracking:</i> [11, 25]	+	+	-
<i>Explainability:</i> [7, 12]	+	-	-
<i>Transparency:</i> [6, 20]	++	+	+
<i>Granularity:</i> [27, 10, 19]	++	+	-

Table 1: Related works. “+” signs indicate wider availability and maturity.

6 Conclusions & Roadmap

The enforcement of data privacy, especially for sensitive information in the health domain, is nowadays legally binding, thanks to current regulations such as the GDPR. Applications and systems dealing with personal data are obliged to follow these directives, even more so in consideration of the broader availability of wearable and sensing devices that collect data about individuals and can potentially make it available for different purposes. Adopting a different perspective, as opposed to top-down approaches for data privacy compliance, in this paper, we introduced a vision for decentralized personal data privacy interactions. This approach is founded on the principles of multi-agent systems, which introduce autonomy, decentralization, and negotiation as essential aspects that allow the establishment of interactions among independent agents, even if they have different goals related to privacy and the use of data. Moreover, we have introduced the use of semantic data models, and, in particular, the DPV ontology, to enable heterogeneous agents to specify privacy policies and consent. Regarding potential threats, differential privacy or techniques related to k-anonymity or

l-diversity can be used to model adversary agents for which privacy agents can progressively develop protection strategies.

This abstract agent model for data privacy handling introduces an overview of how agents can establish relationships and negotiation activities related to data privacy, even though it leaves the question of implementation and deployment of such a system for future work. The challenge of taking this vision to a deployable solution has not to be underestimated. We foresee the following research opportunities in future works:

- (i) The design of domain-specific **vocabularies/ontologies** that describe detailed data processing conditions, purposes and data handling policies;
- (ii) The development of **multi-agent environments** that implement the interactions described above, deployable in mobile and sensing devices.
- (iii) The study and implementation of agent **negotiation protocols** that automate the personal data privacy workflows, such as consent updates, compliance to user preferences, etc.;
- (iv) The specification and validation of **consent and policies** for data privacy, checking automatically for compliance with regulations;
- (v) The **validation and evaluation** of the proposed model, in a real environment and including the verification of strict legal compliance (GDPR).

References

1. Biskup, J., Kern-Isberner, G., Thimm, M.: Towards enforcement of confidentiality in agent interactions. In: Proceedings of the 12th International Workshop on Non-Monotonic Reasoning (NMR'08). pp. 104–112 (2008)
2. Bourgeois, J., Kortuem, G., Kawsar, F.: Trusted and gdpr-compliant research with the internet of things. In: Proceedings of the 8th International Conference on the Internet of Things. pp. 1–8 (2018)
3. Bruschi, D.: Information privacy: Not just gdpr. *Computer Ethics-Philosophical Enquiry (CEPE) Proceedings* **2019**(1), 9 (2019)
4. Buonocunto, P., Giantomassi, A., Marinoni, M., Calvaresi, D., Buttazzo, G.: A limb tracking platform for tele-rehabilitation. *ACM Transactions on Cyber-Physical Systems* **2**(4), 1–23 (2018)
5. Calvaresi, D., Calbimonte, J.P.: Real-time compliant stream processing agents for physical rehabilitation. *Sensors* **20**(3), 746 (2020)
6. Calvaresi, D., Dubovitskaya, A., Retaggi, D., Dragoni, A.F., Schumacher, M.: Trusted registration, negotiation, and service evaluation in multi-agent systems throughout the blockchain technology. In: WI 2018. pp. 56–63. IEEE (2018)
7. Calvaresi, D., Mualla, Y., Najjar, A., Galland, S., Schumacher, M.: Explainable multi-agent systems through blockchain technology. In: Extraamas. pp. 41–58 (2019)
8. Ganti, R.K., Ye, F., Lei, H.: Mobile crowdsensing: current state and future challenges. *IEEE communications Magazine* **49**(11), 32–39 (2011)
9. Jutla, D., Xu, L.: Privacy agents and ontology for the semantic web. *AMCIS 2004 Proceedings* p. 210 (2004)
10. Kanaan, H., Mahmood, K., Sathyan, V.: An ontological model for privacy in emerging decentralized healthcare systems. In: 2017 IEEE 13th International Symposium on Autonomous Decentralized System (ISADS). pp. 107–113. IEEE (2017)

11. Kifor, T., Varga, L.Z., Vazquez-Salceda, J., Alvarez, S., Willmott, S., Miles, S., Moreau, L.: Provenance in agent-mediated healthcare systems. *IEEE Intelligent Systems* **21**(6), 38–46 (2006)
12. Kraus, S., Azaria, A., Fiosina, J., Greve, M., Hazon, N., Kolbe, L., Lembcke, T.B., Müller, J.P., Schleibaum, S., Vollrath, M.: Ai for explaining decisions in multi-agent environments. arXiv preprint arXiv:1910.04404 (2019)
13. Krupa, Y., Vercouter, L.: Handling privacy as contextual integrity in decentralized virtual communities: The privacias framework. *Web Intelligence and Agent Systems: An International Journal* **10**(1), 105–116 (2012)
14. Léauté, T., Faltings, B.: Privacy-preserving multi-agent constraint satisfaction. In: *Intl. Conf. on Computational Science and Engineering*. vol. 3, pp. 17–25 (2009)
15. Miao, C., Jiang, W., Su, L., Li, Y., Guo, S., Qin, Z., Xiao, H., Gao, J., Ren, K.: Cloud-enabled privacy-preserving truth discovery in crowd sensing systems. In: *Proc. ACM Conf. on Embedded Networked Sensor Systems*. pp. 183–196 (2015)
16. Mivule, K., Josyula, D., Turner, C.: An overview of data privacy in multi-agent learning systems. In: *The Fifth International Conference on Advanced Cognitive Technologies and Applications*. pp. 14–20 (2013)
17. Moraffah, B., Sankar, L.: Privacy-guaranteed two-agent interactions using information-theoretic mechanisms. *IEEE Transactions on Information Forensics and Security* **12**(9), 2168–2183 (2017)
18. Palmirani, M., Martoni, M., Rossi, A., Bartolini, C., Robaldo, L.: Pronto: Privacy ontology for legal compliance. In: *Proc. 18th Eur. Conf. Digital Government (ECDG)*. pp. 142–151 (2018)
19. Pandit, H.J., Polleres, A., Bos, B., Brennan, R., Bruegger, B., Ekaputra, F.J., Fernández, J.D., Hamed, R.G., Kiesling, E., Lizar, M., et al.: Creating a vocabulary for data privacy. In: *OTM to Meaningful Internet Systems*. pp. 714–730 (2019)
20. Piolle, G., Demazeau, Y., Caelen, J.: Privacy management in user-centred multi-agent systems. In: *International Workshop on Engineering Societies in the Agents World*. pp. 354–367. Springer (2006)
21. Sanchez, O.R., Torre, I., Knijnenburg, B.P.: Semantic-based privacy settings negotiation and management. *Future Generation Computer Systems* (2019)
22. Sannon, S., Stoll, B., DiFranzo, D., Jung, M.F., Bazarova, N.N.: “i just shared your responses” extending communication privacy management theory to interactions with conversational agents. *Proc. ACM on HCI* **4**, 1–18 (2020)
23. Such, J.M., Espinosa, A., García-Fornes, A.: A survey of privacy in multi-agent systems. *The Knowledge Engineering Review* **29**(3), 314–344 (2014)
24. Thangaraj, M., Ponmalar, P.P., Sujatha, G., Anuradha, S.: Agent based semantic internet of things (iot) in smart health care. In: *Proc. Intl. KMO Conference on The changing face of Knowledge Management Impacting Society*. pp. 1–9 (2016)
25. Vázquez-Salceda, J., Alvarez, S., Kifor, T., Varga, L.Z., Miles, S., Moreau, L., Willmott, S.: Eu provenance project: an open provenance architecture for distributed applications. In: *Agent Technology and e-Health*, pp. 45–63. Springer (2007)
26. Voigt, P., Von dem Bussche, A.: *The eu general data protection regulation (gdpr). A Practical Guide*, 1st Ed., Cham: Springer International Publishing (2017)
27. Wimmer, H., Yoon, V.Y., Sugumaran, V.: A multi-agent system to support evidence based medicine and clinical decision making via data sharing and data privacy. *Decision Support Systems* **88**, 51–66 (2016)
28. Yee, G., Korba, L.: An agent architecture for e-services privacy policy compliance. In: *19th International Conference on Advanced Information Networking and Applications (AINA'05) Volume 1 (AINA papers)*. vol. 1, pp. 374–379. IEEE (2005)