

# REPHRAIN

Protecting citizens online



## Contextual Integrity for Argumentation-based Privacy Reasoning

Gideon Ogunniye, University College London

Nadin Kökciyan, University of Edinburgh

August 2023



# Contextual Integrity for Argumentation-based Privacy Reasoning

Gideon Ogunniye

Department of Science, Technology, Engineering  
and Public Policy (UCL STEaPP),  
University College London,  
London, United Kingdom  
g.ogunniye@ucl.ac.uk

Nadin Kökciyan

School of Informatics,  
University of Edinburgh,  
Edinburgh, United Kingdom  
nadin.kokciyan@ed.ac.uk

## ABSTRACT

Privacy management in online systems is a complex task. Recently, contextual integrity theory has been introduced to model privacy, which considers the social contexts of users before making privacy decisions. However, having a practical application based on this theory is not straightforward. In this paper, we propose an agent-based framework for privacy policy reasoning that combines the power of ontologies together with argumentation techniques to resolve privacy conflicts. First, we propose an ontology that represents the contextual integrity theory. We then introduce an argumentation-based dialogue framework that could: (i) reason about contextual norms to resolve privacy conflicts among agents, and (ii) provide justifications to the agents during multi-party dialogues. We apply our approach to privacy scenarios in various contexts where each scenario has different challenges to address. We conclude with theoretical results to show the effectiveness of the framework.

## KEYWORDS

Privacy; Contextual integrity; Privacy ontology; Argumentation

### ACM Reference Format:

Gideon Ogunniye and Nadin Kökciyan. 2023. Contextual Integrity for Argumentation-based Privacy Reasoning. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 9 pages.

## 1 INTRODUCTION

Online systems are used by billions of users who interact with each other, share data about themselves and also other people. Recent literature shows that online social network (OSN) users share personal information that could lead to privacy breaches [13, 17, 35]. On the other hand, OSNs have become one of the most disruptive communication platforms with high socioeconomic value [22]. For example, sharing content with an unintended audience could result in discriminatory practices [13]. The situation becomes even worse in the IoT domain, where users face the challenge of making privacy decisions in unseen situations [20]. Clearly, online systems have become platforms where users are vulnerable; and there is a need for tools that could help the users to manage their privacy.

Privacy is handled differently in various systems. In general, each piece of content to be shared is associated with a privacy policy that dictates who can access the content in question. One

can provide a specific audience for a post (e.g., Facebook); or the content can be shared publicly or privately (e.g., Twitter). Managing privacy becomes more difficult when multiple people are involved in a content (e.g., a picture representing a group of people) [7, 8, 15]. In current OSNs, users do not have any control about setting multi-party policies; and they can only define a privacy policy of a content they are willing to share themselves (i.e., the content owners). Sharing a multi-party content without consulting the parties involved will most likely result in a multi-party privacy conflict (MPPC) [15, 25]. Furthermore, privacy is very subjective; agreeing on a common privacy policy is not trivial and requires automatic mechanisms to facilitate such negotiations to resolve MPPCs. Kökciyan, Yaglikci and Yolum propose a framework where agents, each representing an OSN user, agree on a sharing decision by using assumption-based argumentation [15]. Such and Rovatsos propose an approach in which each user could assign different degrees of privacy to set the audience of a content [36]. However, privacy is not only about access control. Unlike previous work, we develop a new agent framework in which privacy could be defined individually based on a well-formed theory, namely *Contextual Integrity*. This could help us develop personal privacy assistants (e.g., agents) to assist users in their decision-making based on specific contexts.

In *Contextual integrity (CI)*, privacy is shaped and modified by individual, social, and cultural expectations and norms [26]. CI shows that the information flow is governed by context-dependent norms. These norms are characterized by general institutional and social circumstances; the actors involved and their roles; the information being collected, processed, or shared; and the expected transmission principles [32]. The context of the user defines what is appropriate to share. For example, doctors are expected to keep information about their patients confidential. In an emergency, it would be appropriate to share patient information with other medical professionals, even if explicit consent was not provided. Previous works [2, 33] have used the theory of CI to represent and analyse privacy policies; however, there has been no formal method to represent and reason about the dynamism of contextual norms and to resolve the inconsistencies that arise in privacy policies. Moreover, handling this problem from a multi-party perspective requires new techniques to reason about privacy conflicts.

In this paper, we introduce an agent-based framework for privacy policy reasoning. Our framework is unique in the sense that it combines ontological reasoning with argumentation-based dialogues to make contextual integrity theory practical. For this, we define a new ontology, *PROCI*, to formally represent the privacy domain.

*Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Each agent, who can act on behalf of a user, is equipped with a *PROCI* ontology instance to represent norms for social contexts, privacy preferences, and privacy policies by using semantic rules to make privacy decisions. In addition, agents can exchange justifications for arguments put forward while assisting the user or communicating with other agents through dialogues to resolve privacy conflicts. We illustrate our framework with a couple of practical examples in different social contexts and conclude with theoretical results to show the effectiveness of the proposed approach.

## 2 RELATED WORK

Various approaches are based on agent-based negotiation to agree on a privacy policy in OSNs [12, 25, 37]. Mosca and Such introduce an explainable personal assistant that enables collaboration between agents to identify the optimal privacy policy for collectively owned content [25]. Unlike our approach, the work proposed in [25] does not consider argumentation-based dialogue between agents to exchange information and explanations. Calvaresi, Schumacher and Calbimonte propose an agent-based model in which personal data providers and consumers are embedded in privacy-aware agents, which can negotiate and coordinate data reuse, content, and privacy policies using semantic resources [4]. In these approaches, agents do not have the ability to argue with each other to resolve privacy conflicts.

Argumentation-based approaches, on the other hand, enable the exchange of arguments among agents, where each agent aims to persuade other agents to accept its claims. In recent research [7, 15, 31], argumentation has been shown to be a promising approach to manage privacy in OSNs. In [15], Kokciyan, Yaglikci and Yolum propose an assumption-based argumentation model, where agents represent the users in a social network to argue and decide whether a content should be shared or not. To generate arguments, agents make use of semantic rules that represent their users' privacy constraints. However, these semantic rules do not capture the different social contexts (e.g., friendship contexts) of the users and the contextual factors that may affect users' privacy expectations. In [31], an argumentation-based approach is used to automatically generate explanations and prevent privacy violations in OSNs where contextual factors are not modelled. In our approach, we define a formal language as an ontology to implement contextual integrity theory, we also propose a framework where agents conduct dialogues to prevent privacy violations via argumentation. Oren et al. [29] provide an overview of the body of work on 'arguing about norms'. Similarly, one of the key components of our framework is using argumentation to resolve conflicts between privacy norms.

## 3 BACKGROUND

We first introduce the CI theory. Our proposed framework consists of two main components to implement this: an ontology to formalize CI and an argumentation-based framework to be used by agents for multi-party privacy conflict resolution. We briefly discuss these two components before introducing our framework in Section 4.

### 3.1 Privacy as Contextual Integrity

The theory of contextual integrity (CI) is particularly useful for understanding privacy expectations and the norms of information

transmission in a given context [27]. Nissenbaum suggests that information should be distributed and protected according to the norms governing different social contexts.

In CI, four key components are proposed to formalize privacy: (i) *Contexts* are the situations in which the information flows occur; (ii) *Actors* include senders of information, recipients of information, and information subject (whom the information is about); (iii) *Attributes* are defined as the types of information in the information flow; and (iv) *Transmission Principles* (such as consent, sell, with notice, with a warrant, with authorisation and so on) are the constraints to the information flow from one party to another in a given context. Nissenbaum suggests that contextual integrity is maintained when two types of norms are upheld: *Norms of appropriateness* and *Norms of dissemination*. Norms of appropriateness dictate the type of information about an individual that is appropriate to be revealed in a particular context. Norms of dissemination govern the flow of a third party's personal information from one user to the other.

Representing contexts is a challenging task, and there are various formulations in the literature. Kokciyan and Yolum [18, 20] represent a context as a collection of privacy policies that are semantically similar to each other, where similarity is based on the textual content of the privacy policies. Abowd et al. [1] define context as "any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and the application themselves." Barth et al. [3] define context as the concept that captures the idea that people act and transact in a society not simply as individuals, but as individuals in certain capacities (roles) in distinctive social contexts, such as healthcare, education, etc.

### 3.2 A Formal Language to Implement CI

In our approach, we formalize privacy contexts based on CI theory using a formal language defined as an ontology. An ontology is a way to represent knowledge in a specific domain [23]. We use Web Ontology Language (OWL) to model an ontology for privacy contexts. OWL is characterized by its ability to construct rapid data modeling and enable automatic reasoning [9] and it is more expressive than other ontology languages such as RDF [11]. OWL ontologies consist of classes (i.e., domain concepts) and properties (the relationships between these concepts). They describe instances that are individuals belonging to classes. The ontology captures the semantic rules for the norms that govern information flow in a social context. However, we argue that these norms can be subjective, and their justification may need to be explained or argued about to resolve conflicts. For example, assume that an agent tries to represent a norm regarding when a doctor could share sensitive information about a patient with other doctors without the explicit consent of the patient. The representation of this norm, together with its justification, could vary widely based on cultural differences.

An important step towards ensuring contextual integrity is to allow for communication between agents (human or artificial). These agents must communicate to resolve differences and conflicts of opinions or simply inform each other of pertinent facts about contexts and the impact of changing contexts on privacy policies,

hence the need for argumentation-based dialogues. In the following, we describe ASPIC+ [24, 30] and Dung’s abstract argumentation theory [6] that provide the technical foundations for our proposed argumentation-based framework.

### 3.3 Argumentation Frameworks

In Dung’s abstract argumentation frameworks (AFs) [6], an argument is considered an atomic entity and hence only its interactions with other arguments are modelled. In these frameworks, the internal structure of the arguments and attacks between them is not specified. The emphasis is on the evaluation of the interactions between the arguments to reach a conclusion. [30]. Definition 3.1 provides a formal description of an AF, which is basically a directed graph.

**Definition 3.1.** (Dung’s Argumentation Frameworks (AFs)) [6].  $\mathcal{F} = \langle \mathcal{A}, \mathcal{D} \rangle$  is an AF if  $\mathcal{A}$  is a set of arguments and  $\mathcal{D} \subseteq \mathcal{A} \times \mathcal{A}$  is a binary defeat relation<sup>1</sup> over  $\mathcal{A}$ . Let  $\mathcal{E} \subseteq \mathcal{A}$ :

- $\mathcal{E}$  is *conflict-free* iff there exists no  $\phi_1, \phi_2 \in \mathcal{E}$  such that  $(\phi_1, \phi_2) \in \mathcal{D}$ . Let  $\phi_1 \in \mathcal{E}$ .  $\mathcal{E}$  *defends*  $\phi_1$  iff for every  $(\phi_2, \phi_1) \in \mathcal{D}$ , there exists a  $\phi_3 \in \mathcal{E}$  such that  $(\phi_3, \phi_2) \in \mathcal{D}$ .  $\mathcal{E}$  is an *admissible set* iff  $\mathcal{E}$  is conflict free and defends all its elements.  $\mathcal{E}$  is a *complete extension* iff there are no other elements which it defends.  $\mathcal{E}$  is a *preferred extension* iff it is maximal (with respect to set inclusion) complete extension.

AFs offer semantics for the evaluation of arguments to compute extensions. Here, we focus on preferred semantics, which admits multiple extensions. A preferred extension represents a potentially justified view (which may conflict with other views). If an argument is present in all extensions, then it is *sceptically* justified; while if it is present in at least one extension, it is *credulously* justified.

**3.3.1 Structured Argumentation.** We now define the arguments semantically so that the agents could discuss their different points of view. A general framework for giving structure to arguments is the ASPIC+ framework [24, 30]. ASPIC+ defines an *argumentation system* built from a logical language  $\mathcal{L}$  and defines arguments as inference trees formed by applying strict or defeasible rules to premises that are well-formed formulae (*wff*) in  $\mathcal{L}$ . A *strict rule* means that if one accepts the antecedents, then one must accept the consequent no matter what. A *defeasible rule* means that if one accepts all the antecedents, then one must accept the consequent if there is insufficient reason to reject. The notion of attack between arguments means that a certain *wff* is *contrary* or *contradictory* of certain other *wff*. Definition 3.2 captures a formal definition of an argumentation system.

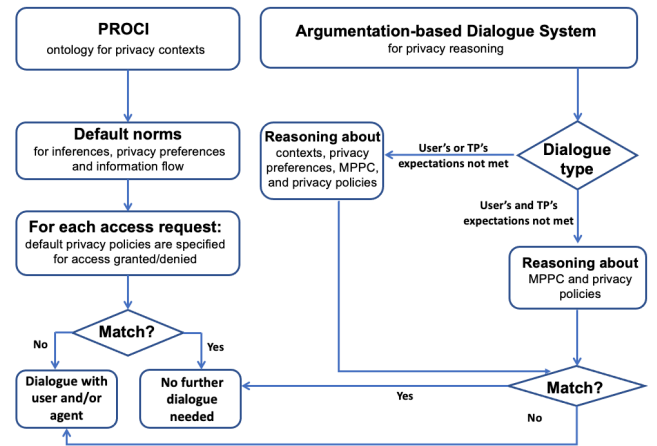
**Definition 3.2.** (Argumentation system). An argumentation system is a triple  $AS = (\mathcal{L}, \mathcal{R}, n)$  where (i)  $\mathcal{L}$  is a logical language closed under negation ( $\neg$ ). (ii)  $\mathcal{R} = \mathcal{R}_s \cup \mathcal{R}_d$  is a set of strict ( $\mathcal{R}_s$ ) and defeasible ( $\mathcal{R}_d$ ) inference rules of the form  $\phi_1, \dots, \phi_n \rightarrow \phi$  and  $\phi_1, \dots, \phi_n \Rightarrow \phi$  respectively (where  $\phi_i, \phi$  are meta-variables ranging over the *wff* in  $\mathcal{L}$ ), and such that  $\mathcal{R}_d \cap \mathcal{R}_s = \emptyset$ . (iii)  $n : \mathcal{R}_d \rightarrow \mathcal{L}$  is a naming convention for defeasible rules.  $n(r)$  is a *wff* in  $\mathcal{L}$  which says that the defeasible rule  $r$  in  $\mathcal{R}$  is applicable.

<sup>1</sup>a *defeat* relation is a version of the *attack* relation indicating a successful attack relation

Arguments in ASPIC+ are constructed from a knowledge base  $\mathcal{K}$ , which contains two disjoint kinds of formulae: *axioms*  $\mathcal{K}_n$  and *ordinary premises*  $\mathcal{K}_p$ . This distinction is important since the agents could attack arguments derived from assumptions. An argumentation theory is a tuple  $AT = (AS, \mathcal{K})$  where  $AS$  is an argumentation system and  $\mathcal{K}$  is a knowledge base in  $AS$ . The formal definition of an argument and an attack is provided in Definition 3.3.

**Definition 3.3.** (Argument and Attack) [24]. An argument  $A$  based on an argumentation theory  $(AS, \mathcal{K})$  with knowledge base  $\mathcal{K}$  and an argumentation system  $(\mathcal{L}, \mathcal{R}, n)$  is: (i)  $\phi$  if  $\phi \in \mathcal{K}$  with:  $Prem(A) = \{\phi\}$ ;  $Conc(A) = \{\phi\}$ ;  $Sub(A) = \{\phi\}$ ;  $TopRule(A) = \text{undefined}$ . (ii)  $A_1, \dots, A_n \rightarrow \mid \Rightarrow \psi$  if  $A_1, \dots, A_n$  are arguments such that there exists a strict or a defeasible rule  $Conc(A_1), \dots, Conc(A_n) \rightarrow \mid \Rightarrow \psi$  in  $\mathcal{R}_s | \mathcal{R}_d$  with  $Prem(A) = Prem(A_1) \cup \dots \cup Prem(A_n)$ ,  $Conc(A) = \{\psi\}$ ,  $Sub(A) = Sub(A_1) \cup \dots \cup Sub(A_n) \cup \{A\}$ ,  $TopRule(A) = Conc(A_1), \dots, Conc(A_n) \rightarrow \mid \Rightarrow \psi$ . (iii)  $A$  attacks  $B$  iff  $A$  undercuts, rebuts, or undermines  $B$ , where  $A$  undercuts  $B$  (on  $B'$ ) iff  $Conc(A) = \neg n(r)$  for some  $B' \in Sub(B)$  such that  $B'$ ’s top rule  $r$  is defeasible.  $A$  rebuts  $B$  (on  $B'$ ) iff  $Conc(A) = \neg \phi$  for some  $B' \in Sub(B)$  of the form  $B'_1, \dots, B'_n \Rightarrow \phi$ .  $A$  undermines  $B$  (on  $\phi$ ) iff  $Conc(A) = \neg \phi$  for an ordinary premise  $\phi$  of  $B$ . Note that  $Prem$  returns all the formulas of  $\mathcal{K}$  called *premises* used to build the argument,  $Conc$  returns its conclusion,  $Sub$  returns all its sub-arguments and  $TopRule$  returns the last inference rule used in the argument.

ASPIC+ defines a set of arguments with binary relation of defeat, that is, it defines argumentation frameworks in the sense of AFs [6], thus making the semantics of abstract argumentation applicable to ASPIC+ (Definition 3.1).



**Figure 1: Agent-Based Framework for Privacy Policy Reasoning, where agent reasoning is based on the PROCI ontology and the argumentation-based dialogue system.**

## 4 AGENT-BASED FRAMEWORK FOR PRIVACY POLICY REASONING

In this section, we describe our proposed framework consisting of two main components: (i) A *privacy ontology based on contextual integrity* (PROCI), which formally specifies the rules that agents

must take into account when taking on various roles in contexts and requesting access to (or granting access to) information containing attributes of the users they represent. (ii) *An argumentation-based dialogue framework* that addresses communication between agents and reasoning about privacy policies. The framework uses the syntax and semantics of the argumentation frameworks presented in Section 3.3 to represent and evaluate privacy dialogues.

Figure 1 depicts our agent-based framework for privacy policy reasoning. The left part of the figure depicts the workflow of a privacy ontology that aims to generate privacy decisions based on some specified norms. The argumentation-based dialogue system manages the different types of dialogue that are possible to initialize. This framework can be instantiated as an interactive dialogue process in each instance of a recommendation for a dialogue.

#### 4.1 Privacy Ontology based on Contextual Integrity (*PROCI*)

We develop an ontology called *PROCI*<sup>2</sup> for contextual integrity. The ontology consists of six distinctive but related classes as depicted in Figure 2: (i) *Agent*; represents a set of agents (i.e., both human and software agents) and their associated instances and attributes, (ii) *Role*; represents a set of all roles that are presently associated with an agent, (iii) *Context*; represents a set of disjoint social contexts (e.g., professional and emergency contexts) that describe the agents, (iv) *Information*; represents a set of facts received or learned about an agent (e.g., an agent’s location or educational qualification, etc.) (v) *Situation*; represents a set of situational conditions or activities of an agent (e.g., an emergency situation, attending a class, presenting a seminar, etc.), and (vi) *Organisation*; represents a set of social or professional organisations that an agent is part of (e.g., a professional association, a religious group, a company, etc.). The ontology also defines a set of properties and relationships that are associated with these classes. In figure 2, note that solid lines represent properties while dashed lines represent subclasses. Following the theory of contextual integrity (CI) as discussed in Section 3.1, *PROCI* includes the following components:

**4.1.1 Contexts, Actors, Roles.** Classes such as *PRContext*, *FRContext* and *EMContext* are used to describe social contexts (professional, friendship, and emergency) of online users, and each of them is a subclass of *Context*. Within a social context, the class *Agent* describes the actors that are involved. We make a distinction between human and software agents since dialogues could be initiated with human agents as well. This class has four subclasses: *User*; is a set of human agents (natural person) who own information or who can be identified directly or indirectly by the information, *ThirdParty*; is a set of human agents that request access to user information, *UserAgent*; the class of software agents that represent users, and *ThirdPartyAgent*; the class of software agents represent third parties. Each *Agent* has a role in a specific context; and some instances are modelled as *:admin-officer*, *:line-manager*, *:em-responder* *:health-and-safety-manager*. To connect agents with roles in a social context, *hasRole* links *Agents* to *Roles*. We define properties such as *isColleagueOf* and *isFriendOf* to represent the relationships

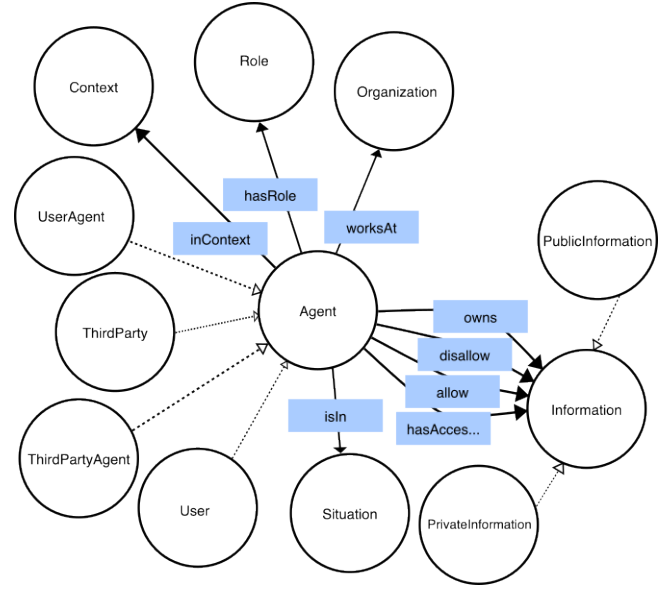


Figure 2: A view from the *PROCI* ontology.

between agents. The following example represents a privacy scenario in an OSN that could be represented by the instantiation of *PROCI*.

**EXAMPLE 1.** *Alice, Bob, and Charlie are all colleagues of each other. Bob prefers to share his location information with Alice and Charlie in a professional context and does not want to reveal this information in a friendship context. Alice is the emergency responder and also the health and safety manager; whereas Charlie is the administrative officer.*

Assume that we model a personal privacy assistant for Bob in this example. An instance of *User* will be instantiated as *:bob*, while *:alice* and *:charlie* are instances of *ThirdParty*. Moreover, user/third party will also have an instance for their software agents. *:bob-a* is an instance of *UserAgent* while *:alice-a* and *:charlie-a* are instances of *ThirdPartyAgent*, respectively. Bob owns some pieces of information such as *:bob-loc* and *:bob-mobile*. Other agents can make *AccessRequests* to access to such data, *bob-a* will be the agent making privacy decisions.

An *Agent* can be in a context with other agents, *inContext* is used to specify this; and the agents can be in multiple contexts [5, 20]. For example, *:bob*, *:alice* and *:charlie* can be in a *PRContext* and an *EMContext* at the same time, and in such combined contexts, it may be appropriate for *:alice* with the role *:health-and-safety-manager* to have access to *:bob*’s mobile information (*:bob-mobile*) but not *:charlie* with the role *:admin-officer*.

**4.1.2 Norms of Appropriateness/Dissemination.** We define three types of norms to govern the information flow in a social context, as introduced in Section 3.1: Access norms (A), Inference norms (I) and Privacy norms (P). The access norms generalize the *norms of appropriateness* and *norms of dissemination* of contextual integrity. These norms specify whether access requests are allowed or denied; the *allow* and *disallow* predicates appear in the head of access norms. An access norm A allows a third party to access private

<sup>2</sup>The ontology is shared as part of our supplementary material. Note that *PROCI* is populated with a representative list of classes and relations.

information of a user *if* its conditions are satisfied; otherwise, the access is denied. Inference norms enable a software agent (a user agent or a third-party agent) to derive new information from existing knowledge in its ontology. In a privacy norm, the user specifies which type of access request should be granted or rejected.

In this work, we assume that the software agents are aware of the norms of their users. Some related work focus on learning such norms [16, 21] which is out of our scope. Semantic Web Rule Language (SWRL) [10] is used to represent the CI norms. The norms are defined as a set of precedent and consequent states consisting of the conjunction of atoms (i.e.,  $\text{Body} \rightarrow \text{Head}$ ), which means that if the body holds, then the head must hold. Here, the atoms are of the form  $c(x)$  and  $P(x, y)$  where  $c$  is a class name (e.g., *Information*) and  $P$  is a property name (e.g., *hasRole*), which are defined in the ontology.  $x$  and  $y$  are variables prefixed with a question mark (e.g. *?user*), instance names (e.g. *:bob*) or literals (e.g. *:true*).

In Table 1, we show an example set of norms of inference (I), privacy preference (P) and access request (A). Each norm is modelled as a SWRL rule. For example, Bob’s agent uses the three inference norms  $I_1$ ,  $I_2$ , and  $I_3$  to infer whether a user and a third party are in a friendship, professional, or emergency context, respectively. For example, the norm  $I_3$  states that if there is an emergency situation and a user (*?user*) is in the emergency situation; and there is a third party (*?tp*) with a role of emergency responder; then *?user* and *?tp* are in an emergency context. *:bob* has two privacy norms  $P_1$  and  $P_2$  that state the contexts in which he prefers to share (resp., not share) his location information with a third party. Access norms  $A_1$ ,  $A_2$  and  $A_3$  are used by Bob’s agent *:bob-a* to allow or deny access to Bob’s location information for access requests from Alice (a third party). For instance, norm  $A_3$  states that if there is an access request from a third party (*:alice*) for Bob’s location information *:bob-loc* (that is, *hasAR(:alice, :bob-loc)*), the access request should be denied in a friendship context *inFRContext(:bob, :alice)*.

**4.1.3 Privacy Conflicts.** Access rules represent CI-aware norms that the user may not be fully aware of. In some cases, access norms and privacy norms can result in a similar privacy decision. For example, in Table 1, the access norm  $A_3$  is in agreement with Bob’s privacy rule  $P_2$  in the sense that they lead to the same outcome. However, agents can infer additional information that could lead to some disagreement. Therefore, the preference and access norms could be in conflict. Definition 4.1 captures this.

**Definition 4.1 (Privacy Conflict).** For a specific user  $u$ , given a social context with an access request for user information from a third party, a privacy conflict occurs *iff*  $\text{Head}(A_i) \neq \text{Head}(P_j)$ , where  $A_i$  is an access norm of  $u$ ,  $P_j$  is a privacy norm of  $u$  and  $\forall i, j \in \mathbb{N}$ .

*PROCI* explicitly represents privacy conflicts to allow agents to reason about them. As depicted in Figure 1, the agent needs to initialize a dialogue when there is a mismatch among norms (that is, a privacy conflict). We now introduce an argumentation-based dialogue framework for agents to resolve such privacy conflicts.

## 4.2 Argumentation-based Dialogue Framework

There are instances where a user agent needs to negotiate with a specific user about whether or not a private information of the user should be shared (*reasoning about user privacy preferences*) or the

**Table 1: Some Norms for Bob as SWRL Rules**

$I_1$ : <i>isFriendOf</i> ( <i>?user, ?tp</i> ) $\rightarrow$ <i>inFRContext</i> ( <i>?user, ?tp</i> )
$I_2$ : <i>workAt</i> ( <i>?user, ?office</i> ), <i>workAt</i> ( <i>?tp, ?office</i> ), <i>hasRole</i> ( <i>?tp, ?role</i> ), <i>isColleagueOf</i> ( <i>?user, ?tp</i> ) $\rightarrow$ <i>inPRContext</i> ( <i>?user, ?tp</i> )
$I_3$ : <i>Emergency</i> ( <i>?em</i> ), <i>isInEmergency</i> ( <i>?user, ?em</i> ), <i>hasRole</i> ( <i>?tp, :em-responder</i> ) $\rightarrow$ <i>inEMContext</i> ( <i>?user, ?tp</i> )
$P_1$ : <i>owns</i> ( <i>:bob, :bob-loc</i> ), <i>inPRContext</i> ( <i>:bob, ?tp</i> ) $\rightarrow$ <i>allow</i> ( <i>?tp, :bob-loc</i> )
$P_2$ : <i>owns</i> ( <i>:bob, :bob-loc</i> ), <i>inFRContext</i> ( <i>:bob, ?tp</i> ) $\rightarrow$ <i>disallow</i> ( <i>?tp, :bob-loc</i> )
$A_1$ : <i>owns</i> ( <i>:bob, :bob-loc</i> ), <i>inPRContext</i> ( <i>:bob, :alice</i> ), <i>hasAR</i> ( <i>:alice, :bob-loc</i> ) $\rightarrow$ <i>allow</i> ( <i>:alice, :bob-loc</i> )
$A_2$ : <i>owns</i> ( <i>:bob, :bob-mobile</i> ), <i>inEMContext</i> ( <i>:bob, :alice</i> ), <i>hasAR</i> ( <i>:alice, :bob-mobile</i> ) $\rightarrow$ <i>allow</i> ( <i>:alice, :bob-mobile</i> )
$A_3$ : <i>owns</i> ( <i>:bob, :bob-loc</i> ), <i>inFRContext</i> ( <i>:bob, :alice</i> ), <i>hasAR</i> ( <i>:alice, :bob-loc</i> ) $\rightarrow$ <i>disallow</i> ( <i>:alice, :bob-loc</i> )

context in which the user can share a private information (*reasoning about contexts*). Similarly, a user agent may need to negotiate with a third party agent to resolve multiparty privacy conflict (*reasoning about MPPCs*). Here, argumentation-based dialogues are conceived as the underlying mechanism by which an agent communicates with a specific user or with a third party to resolve privacy conflicts.

**4.2.1 The Dialogue Protocol.** A dialogue between a user agent and a specific user or between a user agent and a third party agent consists of a sequence of *moves*, where each move references both a statement and the user agent/user who made the statement. A statement can be a request for information, a provided information, or a privacy decision to grant (resp., decline) an access request. This is captured in Definition 4.2.


**Definition 4.2.** A dialogue  $D$  consists of a sequence of iterations such that  $D = [[M_1^1, \dots, M_x^1], \dots, [M_1^t, \dots, M_x^t]]$ ; which involves  $n$  participants  $p_1, \dots, p_n$  where ( $n \geq 2$ ). Within a dialogue  $D$ , iteration  $j$  consists of a sequence of moves  $[M_1^j, \dots, M_x^j]$ .

A user agent evaluates the set of statements exchanged within an iteration to update its privacy knowledge. Within each iteration, there is a claim to be discussed and arguments that attack or defend the claim. Note that a claim is abstractly represented as an argument. An iteration therefore represents a sub-discussion focused around a single topic of the dialogue, which can be treated in an atomic manner.

The dialogue protocol is depicted in Figure 3, which describes the set of legal moves that are permitted in each iteration. Each node represents the type of move an agent can make in a dialogue, and the outgoing arcs of a node indicate possible responding moves. In the protocol, the following moves are defined: *assert*( $Ag_i, \phi, t$ ), *contradict*( $Ag_i, \phi_1, \phi_2, t$ ), *retract*( $Ag_i, \phi, t$ ), *exit* and *update*. Note that  $Ag_i$  represents the user agent or a user/third-party agent who made a move. A user agent uses the *assert*( $Ag_i, \phi, t$ ) move



to start a dialogue by submitting a claim  $\phi$  in iteration  $t$ . The claim could be a context inferred from its ontology that the user it represents or a third party agent needs to be aware of (reasoning about context) or an explanation of a privacy decision in a case where the privacy decision is in conflict with the privacy preference of the user or the third party agent (reasoning about MPPC) or a new information for the user to update its privacy preference (reasoning about user privacy preferences). Similarly, a user agent uses the  $contradict(Ag_i, \phi_1, \phi_2, t)$  move labeled  $M_{x+1}^t$  to make a claim  $\phi_2$  to attack a previous claim  $\phi_1$  made by a user/third party agent in move  $M_x^t$ . The claim  $\phi_2$  of the move  $M_{x+1}^t$  must be *relevant* to the claim  $\phi_1$  of the move  $M_x^t$ . Specifically, a claim  $\phi_2$  is relevant to  $\phi_1$  if it attacks it (cf. Definition 3.1). A user agent or a user/third party agent uses the  $retract(Ag_i, \phi, t)$  move to retract its previous claim. A user agent or a user/third party agent uses the  $exit$  move to exit an iteration. This move is made when the user agent or a user/third party agent has no more  $contradict(Ag_i, \phi_1, \phi_2, t)$  move to advance within the iteration. When an iteration is completed (shown by the terminal node *update* in the figure), the dialogue state is updated. The dialogue then proceeds to the next iteration or may terminate.



**Figure 3: The legal moves of agents from the dialogue protocol**

**4.2.2 Combining PROCI with Dialogues.** We assume that a user agent or a user/third party agent usually has a set of norms/rules to generate its set of arguments in a dialogue. At any point in time, a user agent or a user/third party agent may *add* argument to a dialogue if it is not already present within it or *retract* argument (if it was already present) by adding another argument that indicates retraction of a previous argument. Therefore, a dialogue usually results in a set of arguments and attack or defeat relations arising from conflicts between arguments. In our framework (Figure 1), we apply the ASPIC+ argumentation theory to represent and evaluate a dialogue. Note that we do not consider an ordering  $\leq$  of the elements of  $\mathcal{K}$  [24]. To resolve an attack between arguments to a defeat, an attacker wins unless the attacked argument is defended by other undefeated arguments. We define the mapping from PROCI to our framework in Definition 4.3.

**Definition 4.3.** Given PROCI and  $AS = (\mathcal{L}, \mathcal{R}, n)$ , with knowledge base  $\mathcal{K}$ , in our framework, we adapt ASPIC+ argumentation theory  $AT = (AS, \mathcal{K})$  as follows:

- (1)  $\mathcal{L}$  is the formal language defined in PROCI; arguments and attacks are defined based on Definition 3.3; contrariness is as defined in Section 3.3. For simplicity, we assume that different contexts are contrary of each other.
- (2)  $\mathcal{K} = \mathcal{K}_p \cup \mathcal{K}_a$  where  $\mathcal{K}_{p|a} = \{v \in \mathcal{L} | v \text{ is a property/class axiom in } \mathcal{PROCI}\}$  such that: (i) In PROCI,  $\mathcal{K}_p$  is a set of properties/property axioms. For example, the property  $\text{hasRole}(\text{alice}, \text{admin-officer})$  is an ordinary premise that can be attacked. (ii)  $\mathcal{K}_a$  is a set of class/class expression axioms in PROCI. For example,  $\text{disjointClasses}(\text{User}, \text{ThirdParty})$  is a class axiom from PROCI. Intuitively, arguments cannot be attacked on their axiom premises.
- (3)  $\mathcal{R} = \mathcal{R}_d \cup \mathcal{R}_s$  is the smallest set of inference, privacy, and access rules such that: (i) Inference rules are strict rules set by the user agent as default rules; hence, they are mapped to  $\mathcal{R}_s$ . (ii) Access and Privacy rules in the ontology are subjective perceptions and/or evaluations of the user agent and user/third party agent, respectively, and are therefore mapped to defeasible rules ( $\mathcal{R}_d$ ).

A user agent  $Ag_i$  evaluates the arguments that have been put forward in a dialogue. Definition 4.4 formally defines an agent-specific argumentation framework.

**Definition 4.4 (Agent-specific Argumentation Framework).** Let  $Ag_i$  be a user agent, the argumentation framework of  $Ag_i$  is a tuple  $AF_{Ag_i} = \langle \mathcal{A}_D, \mathcal{D}_D \rangle$  where  $\mathcal{A}_D$  is a set of arguments in a dialogue  $D$  and  $\mathcal{D}_D \subseteq \mathcal{A}_D \times \mathcal{A}_D$  is a defeat relation.

Note that an  $AF_{Ag_i}$  can have multiple extensions when the user agent performs the reasoning under preferred semantics [6]. To define the set of justified arguments in our framework, we borrow the notions of sceptical and credulous acceptability (Definition 4.5).

**Definition 4.5 (Acceptability).** Given  $AF_{Ag_i} = \langle \mathcal{A}_D, \mathcal{D}_D \rangle$ , a set of arguments  $\mathcal{E}_D \subseteq \mathcal{A}_D$  in  $AF_{Ag_i}$  is a preferred extension for a user agent  $Ag_i$  if it is a maximal (with respect to set inclusion) complete extension obtained from  $AF_{Ag_i}$ . An argument  $\phi$  is sceptically justified if it is in every  $\mathcal{E}_D$  obtained from  $AF_{Ag_i}$ . On the other hand,  $\phi$  is credulously justified if it is in some  $\mathcal{E}_D$ .

### 4.3 Dialogue between a User Agent and a User about Privacy Preferences and Contexts

In Table 2, we illustrate a dialogue between a *user* (Bob) and his *agent*. The example shows the conflict between Bob's preference and the privacy decision of his agent regarding a third-party's request for his location information. This example also shows how a dialogue could help agents update their information about privacy contexts. Consider the following example where Alice wants to access Bob's location information.

**EXAMPLE 2.** *There is a gas explosion reported near Bob's residence. Alice is worried about Bob and wants to access his location. This access request results in a dialogue between Bob and his agent bob-a.*

We use the following notation to represent the moves of the agents in a dialogue. In a move  $M_x^y$ ,  $x$  denotes the move index

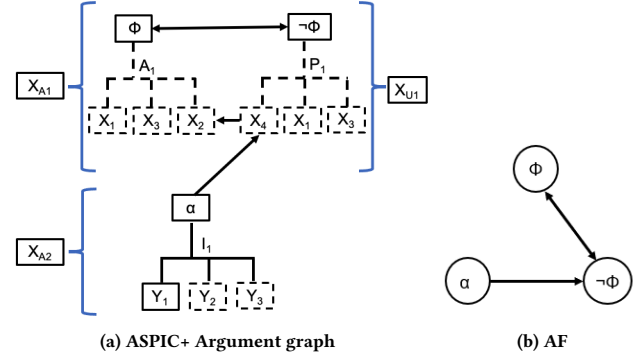
number, while  $y$  denotes the iteration in which the move was made. In Table 2, the move  $M_1^1$  is an *assert*( $Ag_i, \phi, t$ ) move of Bob's agent (bob-a) to start the dialogue. The argument  $\phi$  in the move is generated from the rules inferred from the agent ontology, and suggests that the location information should be shared. The move  $M_2^1$  is a *contradict*( $Ag_j, \phi_1, \phi_2, t$ ) move by Bob, because Bob does not want to share his location information in a friendship context (i.e., on Facebook). bob-a then makes a move to explain to Bob that there was an explosion nearby and therefore the users were in an emergency context. In other words, bob-a makes the move  $M_3^1$  to *contradict*( $Ag_i, \phi_1, \phi_2, t$ ). The argument  $\neg\phi$  in move  $M_2^1$  by Bob can be generated from the rules that are inferred from its privacy settings on OSNs and / or can be taken as input, as we consider here. We map these rules to the adapted ASPIC + argumentation theory (see Table 4.3). Note that the argument in move  $M_1^1$  is rebutted by the argument in move  $M_2^1$ . Furthermore, the rules in moves  $M_1^1$  and  $M_2^1$  are the access and privacy rules of Bob's agent and Bob (that are involved in the dialogue, respectively) and are represented as defeasible rules (c.f., Definition 4.3). The argument in move  $M_3^1$  defeats the argument in  $M_2^1$ , as it is an undefeated argument. The inference rule in the move  $M_3^1$  is represented as a strict rule (c.f., Definition 4.3). Since Bob did not make any further *contradict*( $Ag_j, \phi_1, \phi_2, t$ ) move, Bob's agent  $Ag_i$  made an *exit* move to exit the iteration and an *update* move to update the dialogue state, respectively. Note that we do not represent these two moves in Table 2 as their representation is trivial.

**Table 2: An example dialogue between an agent and its user to agree on a privacy decision**

Moves/Players	SWRL Rules	Arguments
$M_1^1$ by :bob-a :alice would like to access your location on Facebook	$X_1$ : owns(:bob, :bob-loc) $X_2$ : inEMContext(:bob, :alice) $X_3$ : hasAR(:alice, :bob-loc) $\Rightarrow \phi$ : allow(:alice, :bob-loc)	$\phi$ : share location with :alice
$M_2^1$ by :bob Bob wants to apply his privacy setting in a friendship context.	$X_1$ : owns(:bob, :bob-loc) $X_4$ : inFRContext(:bob, :alice) $X_3$ : hasAR(:alice, :bob-loc) $\Rightarrow \neg\phi$ : disallow(:alice, :bob-loc)	$\neg\phi$ : do not share location with :alice
$M_3^1$ by :bob-a A gas explosion has been reported near your residence	$Y_1$ : Emergency(:gas-exp) $Y_2$ : isInEmergency(:bob, :gas-exp) $Y_3$ : hasRole(:alice, :em-responder) $\rightarrow \alpha$ : inEMContext(:bob, :alice)	$\alpha$ : there is an emergency

Arguments in Table 2 and their corresponding defeat relations are abstracted into an ASPIC + argument graph, as shown in Figure 4a where the rectangle with a solid line denotes an axiom  $\{Y_1\}$ , while those with dotted lines denote ordinary premises such as  $X_1$  and  $Y_2$ . Strict rules are denoted with solid lines, whereas defeasible rules are denoted with dotted lines. Rules  $A_x$ ,  $P_x$  and  $I_x$  denote the access, privacy, and inference rules, respectively.  $X_{Ai}$  denotes an

argument moved by Bob's agent, while  $X_{Ui}$  denotes the one moved by Bob. An argument  $X_{A1}$  for  $\phi$  (i.e. with conclusion  $\phi$ ) is shown in Figure 4a with the premises at the bottom and the conclusion at the top of the tree. The directed arrows denote defeat relations. The ASPIC+ graph can be further abstracted into the AF (Figure 4b). The justified arguments of AF are then evaluated using the preferred semantics of AFs as  $\{\phi, \alpha\}$  [6]. In our example, Alice is granted access to Bob's location information :bob-loc as arguments  $\phi$  and  $\alpha$  are sceptically justified.



**Figure 4: (a) depicts the dialogue between two agents (Table 2). (b) is the abstract argument graph used for the evaluation of the winning arguments under preferred semantics.**

#### 4.4 Handling Multi-Party Privacy Conflicts

In this section, we focus on a multi-party dialogue between two agents where a third-party agent (:charlie-a) dialogue with Bob's agent about whether or not to share Bob's data with another third-party (:alice-a). Hence, the agents try to agree on a privacy decision automatically by conducting a dialogue. Note that the complete set of rules can be found in the *PROCI* ontology. Consider the following example where multi-party privacy conflicts occur.

**EXAMPLE 3.** Bob is involved in an accident. alice-a, the agent of Alice, is asking the agent of Charlie (charlie-a) to share Bob's salary information. charlie-a is trying to persuade the agent of Bob (bob-a) to share Bob's salary information with (alice-a).

Charlie may have access to Bob's salary information (as an admin-officer), while Alice is denied such access by the agent of Bob as she is the health-and-safety-manager. The norms of dissemination prevent Charlie from sharing Bob's salary information with Alice without Bob's consent. We represent a sample dialogue in Table 3. The argument of move  $M_1^1$  by :charlie-a claims that Bob's salary information should be shared with a third-party (Alice), because charlie-a has information about Bob being involved in an accident. This piece of information is missing for bob-a at the time of the dialogue; hence, this argument is attacked by the argument in move  $M_2^1$  by Bob's agent :bob-a based on the access rules defined in its ontology. The argument of move  $M_3^1$  by :charlie-a defeats the argument of move  $M_2^1$ . Now that bob-a has information about the accident context, it is appropriate to share Bob's salary information



**Table 3: An example dialogue between two agents (:bob-a and charlie-a) that involves three different users.**

Moves/Players	SWRL Rules	Arguments
$M_1^1$ by :charlie-a Alice should access Bob's salary information	$X_5: \text{owns}(:\text{bob}, :\text{bob-salary})$ $X_6: \text{hasRole}(:\text{alice}, :\text{h-safety-officer})$ $X_7: \text{inAccContext}(:\text{bob}, :\text{alice})$ $X_8: \text{hasAR}(:\text{alice}, :\text{bob-salary})$ $\Rightarrow \psi: \text{allow}(:\text{alice}, :\text{bob-salary})$	$\psi: \text{share}$ Bob's salary information with Alice
$M_2^1$ by :bob-a Alice should not access Bob's salary information	$X_5: \text{owns}(:\text{bob}, :\text{bob-salary})$ $X_6: \text{hasRole}(:\text{alice}, :\text{h-safety-officer})$ $X_9: \text{inPRContext}(:\text{bob}, :\text{alice})$ $X_8: \text{hasAR}(:\text{alice}, :\text{bob-salary})$ $\Rightarrow \neg\psi: \text{disallow}(:\text{alice}, :\text{bob-salary})$	$\neg\psi: \text{do not}$ share Bob's salary information with Alice
$M_3^1$ by :charlie-a The information is needed to compensate for the injury	$X_6: \text{hasRole}(:\text{alice}, :\text{h-safety-officer})$ $Y_4: \text{inAccident}(:\text{bob}, :\text{true})$ $\rightarrow \varphi: \text{inAccContext}(:\text{bob}, :\text{alice})$	$\varphi: \text{Bob was}$ involved in accident

with Alice, and Charlie sharing this information with Alice would not violate the norms of dissemination anymore.

The example in Section 4.3 illustrates how an argumentation-based dialogue can be used as a mechanism to ensure appropriate access to the information of a user (i.e., norms of appropriateness of CI); here, we illustrate the norms of dissemination. It is important to note that this type of dialogue is useful to facilitate exchange of explanations between agents/users over privacy preferences and expectations. The explanations are important to improve the understanding of users about particular privacy situations. We will work on the explanation aspects of our framework as part of our future work. Furthermore, some agents may be malicious or incompetent, and — to achieve desirable dialogical outcomes — the inputs from these agents should be discounted for lack of trust [28].

## 4.5 Theoretical Results

We have implemented the *PROCI* ontology using the Protégé tool<sup>3</sup>, which has automated reasoners (e.g. Pellet) to detect inconsistencies and redundant knowledge. The consistency of the ontology depends on the reasoner being used to apply further inference. However, due to the overlap of social contexts or the hierarchical relationship between actors, their roles, and attributes, it is expected to come up with inconsistencies in the formulation of rules for a privacy policy. For example, there may be rules for both positive and negative authorisation of access to the same piece of information, resulting in privacy conflicts; which should be addressed externally to increase transparency in the reasoning process as we do here by using argumentation. Definition 4.6 defines what a consistent ontology instance is.

**Definition 4.6 (Consistency).** A *PROCI* instance  $\theta$  is consistent for a privacy decision  $\alpha$  if there is no contradiction in the set of access rules  $A_1, \dots, A_n$  and privacy rules  $P_1, \dots, P_n$  in *PROCI* or the set of arguments  $X_1, \dots, X_n$  in *AT* that support  $\alpha$ . That is,  $A_1, \dots, A_n \wedge P_1, \dots, P_n | X_1, \dots, X_n \Rightarrow \alpha$ .

<sup>3</sup><http://protege.stanford.edu/>

**THEOREM 4.7.** *Inference, privacy, and access rules are consistent in the ontology given a particular privacy scenario.*

**PROOF SKETCH.** Each agent is equipped with an ontology that is consistent thanks to the underlying reasoner, which is Pellet [34] in our case. Pellet is sound and complete with respect to OWL. Therefore, each of the rules in the ontology is semantically consistent. However, domain-specific conflicts may arise between different rules, which we address through argumentation-based dialogues. To resolve inconsistencies between conflicting rules, we apply the semantics of the ASPIC+ and Dung argumentation frameworks (see Section 4.3 and Tables 2 and 3).

**THEOREM 4.8.** *For any  $AF_{Ag_i}$  where the agent adopts preferred semantics to evaluate arguments, the set of computed extensions (i.e., the sets of justified arguments) are consistent, complete, and sound. The proposed agent-specific argumentation framework will preserve consistency of a *PROCI* instance to make privacy decisions.*

**PROOF SKETCH.** The preferred semantics ensure that there is an existence of at least one extension. The extensions computed under preferred semantics are conflict-free; therefore, we can say that preferred semantics can be applied to resolve inconsistencies in the rules to obtain a consistent *PROCI* instance considering the agent's preferences to choose among alternative extensions. With respect to completeness and soundness, completeness can be proven for a finite set of arguments in a dialogue. The dialogue protocol described in Section 4.2 defines the dialogue moves allowed of the agents in a dialogue, and the criteria for the termination and evaluation of a dialogue. These protocols ensure that a dialogue results in a finite set of arguments that are automatically derived from the agents' ontology. Therefore, all possible arguments are generated for a particular privacy scenario. Then completeness follows, since every preferred extension is maximal (with respect to set inclusion) complete extension. Likewise, soundness follows, since each of the preferred extensions includes valid arguments derived from the agent's ontology. Hence, each preferred extension consists of a set of admissible arguments to make a privacy decision.

## 5 CONCLUSION

In this paper, we have proposed an agent-based framework to handle privacy policy reasoning in social contexts. Our framework builds on the theory of contextual integrity to represent social contexts, the actors involved, and the norms to govern information flow. We have formally defined a new ontology, *PROCI*, and the argumentation-based dialogue framework. We have presented two examples, where we illustrate how user agent and user/third-party agents apply the inference, privacy, and access rules within their ontologies for privacy policies and how argumentation provides an underlying mechanism for resolving the conflicts. Our theoretical results also provide completeness and soundness guarantees.

As future work, it would be interesting to incorporate machine learning (e.g., reinforcement learning) to learn access rules based on the privacy behaviour of the user in online systems such as OSN and IoT [19]. We also plan to develop a chatbot to manage the dialogue model of our framework, similar to the work done in [14], where the agent could explain the decision-making process to the user. Such an interaction could then be used to choose among alternative privacy decisions while keeping the human in the loop.

## ACKNOWLEDGMENTS

This research was funded by the UKRI Strategic Priorities Fund via the REPHRAIN research centre.

## REFERENCES

- [1] Gregory D Abowd, Anind K Dey, Peter J Brown, Nigel Davies, Mark Smith, and Pete Steggles. 1999. Towards a better understanding of context and context-awareness. In *International symposium on handheld and ubiquitous computing*. Springer, 304–307.
- [2] Noah Apthorpe, Sarah Varghese, and Nick Feamster. 2019. Evaluating the Contextual Integrity of Privacy Regulation: Parents' {IoT} Toy Privacy Norms Versus {COPPA}. In *28th USENIX Security Symposium (USENIX Security 19)*. 123–140.
- [3] Adam Barth, Anupam Datta, John C Mitchell, and Helen Nissenbaum. 2006. Privacy and contextual integrity: Framework and applications. In *2006 IEEE symposium on security and privacy (S&P'06)*. IEEE, 15–pp.
- [4] Davide Calvaresi, Michael Schumacher, and Jean-Paul Calbimonte. 2020. Personal data privacy semantics in multi-agent systems interactions. In *International Conference on Practical Applications of Agents and Multi-Agent Systems*. Springer, 55–67.
- [5] Natalia Criado and Jose M Such. 2015. Implicit contextual integrity in online social networks. *Information Sciences* 325 (2015), 48–69.
- [6] Phan Minh Dung. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial intelligence* 77, 2 (1995), 321–357.
- [7] Ricard L Fogues, Pradeep K Murukanniah, Jose M Such, and Munindar P Singh. 2017. Sharing policies in multiuser privacy scenarios: Incorporating context, preferences, and arguments in decision making. *ACM Transactions on Computer-Human Interaction (TOCHI)* 24, 1 (2017), 1–29.
- [8] Ricard L Fogues, Pradeep Murukanniah, Jose M Such, Agustin Espinosa, Ana Garcia-Fornes, and Munindar Singh. 2015. Argumentation for multi-party privacy management. (2015).
- [9] Jeff Heflin et al. 2007. An introduction to the owl web ontology language. *Lehigh University. National Science Foundation (NSF)* 7 (2007).
- [10] Ian Horrocks, Peter F Patel-Schneider, Harold Boley, Said Tabet, Benjamin Grosz, Mike Dean, et al. 2004. SWRL: A semantic web rule language combining OWL and RuleML. *W3C Member submission* 21, 79 (2004), 1–31.
- [11] Imrul Kayes and Adriana Iamnitchi. 2013. Aegis: A semantic implementation of privacy as contextual integrity in social ecosystems. In *2013 Eleventh Annual Conference on Privacy, Security and Trust*. IEEE, 88–97.
- [12] Dilara Kekulluoglu, Nadin Kökciyan, and Pinar Yolum. 2018. Preserving Privacy As Social Responsibility in Online Social Networks. *ACM Transactions on Internet Technology (TOIT)* 18, 4, Article 42 (2018), 22 pages.
- [13] Dilara Kekulluoglu, Walid Magdy, and Kami Vaniea. 2020. Analysing privacy leakage of life events on twitter. In *12th ACM conference on web science*. 287–294.
- [14] Nadin Kökciyan, Isabel Sassoon, Elizabeth Sklar, Sanjay Modgil, and Simon Parsons. 2021. Applying Metalevel Argumentation Frameworks to Support Medical Decision Making. *IEEE Intelligent Systems* 36, 2 (2021), 64–71.
- [15] Nadin Kökciyan, Nefise Yaglikci, and Pinar Yolum. 2017. An Argumentation Approach for Resolving Privacy Disputes in Online Social Networks. *ACM Transactions on Internet Technology (TOIT)* 17, 3, Article 27 (2017), 27:1–27:22 pages.
- [16] Nadin Kökciyan, Nefise Yaglikci, and Pinar Yolum. 2017. An argumentation approach for resolving privacy disputes in online social networks. *ACM Transactions on Internet Technology (TOIT)* 17, 3 (2017), 1–22.
- [17] Nadin Kökciyan and Pinar Yolum. 2016. PriGuard: A Semantic Approach to Detect Privacy Violations in Online Social Networks. *IEEE Transactions on Knowledge and Data Engineering (TKDE)* 28, 10 (2016), 2724–2737.
- [18] Nadin Kökciyan and Pinar Yolum. 2017. Context-Based Reasoning on Privacy in Internet of Things. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI)*. 4738–4744.
- [19] Nadin Kökciyan and Pinar Yolum. 2020. TURP: Managing Trust for Regulating Privacy in Internet of Things. *IEEE Internet Computing* 24, 6 (2020), 9–16.
- [20] Nadin Kökciyan and Pinar Yolum. 2022. Taking Situation-Based Privacy Decisions: Privacy Assistants Working with Humans. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, Lud De Raedt (Ed.). International Joint Conferences on Artificial Intelligence Organization, 703–709. Main Track.
- [21] A Kurtan and Pinar Yolum. 2021. Assisting humans in privacy management: an agent-based approach. *Autonomous Agents and Multi-Agent Systems* 35, 1 (2021), 1–33.
- [22] Roberto Marmo. 2015. Social Commerce Using Social Network and E-Commerce. In *Encyclopedia of Information Science and Technology, Third Edition*. IGI Global, 2351–2359.
- [23] Deborah L McGuinness, Frank Van Harmelen, et al. 2004. OWL web ontology language overview. *W3C recommendation* 10, 10 (2004), 2004.
- [24] Sanjay Modgil and Henry Prakken. 2014. The ASPIC+ framework for structured argumentation: a tutorial. *Argument & Computation* 5, 1 (2014), 31–62.
- [25] Francesca Mosca and Jose Such. 2022. An explainable assistant for multiuser privacy. *Autonomous Agents and Multi-Agent Systems* 36, 1 (2022), 1–45.
- [26] Helen Nissenbaum. 2009. *Privacy in context*. Stanford University Press.
- [27] Helen Nissenbaum. 2011. A contextual approach to privacy online. *Daedalus* 140, 4 (2011), 32–48.
- [28] Gideon Ogunniye, Alice Toniolo, and Nir Oren. 2017. A Dynamic Model of Trust in Dialogues. In *International Workshop on Theorie and Applications of Formal Argumentation*. Springer, 211–226.
- [29] Nir Oren, Antonino Rotolo, Leendert van der Torre, and Serena Villata. 2013. Norms and argumentation. In *Agreement Technologies*. Springer, 233–249.
- [30] Henry Prakken. 2010. An abstract framework for argumentation with structured arguments. *Argument and Computation* 1, 2 (2010), 93–124.
- [31] Ramon Ruiz-Dolz, José Alemany, Stella Heras, and Ana Garcia-Fornes. 2019. Automatic Generation of Explanations to Prevent Privacy Violations.. In *XAILA@ JURIX*.
- [32] Florian Schaub, Bastian Könings, and Michael Weber. 2015. Context-adaptive privacy: Leveraging context awareness to support privacy decision making. *IEEE Pervasive Computing* 14, 1 (2015), 34–43.
- [33] Yan Shvartzshnaider, Noah Apthorpe, Nick Feamster, and Helen Nissenbaum. 2019. Going against the (appropriate) flow: A contextual integrity approach to privacy policy analysis. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 7. 162–170.
- [34] Evren Sirin, Bijan Parsia, Bernardo Cuenca Grau, Aditya Kalyanpur, and Yarden Katz. 2007. Pellet: A practical OWL-DL reasoner. *Journal of Web Semantics* 5, 2 (2007), 51–53. Software Engineering and the Semantic Web.
- [35] Anna C. Squicciarini, Cornelia Caragea, and Rahul Balakavi. 2017. Toward Automated Online Photo Privacy. *ACM Transactions on the Web* 11, 1, Article 2 (April 2017), 29 pages.
- [36] Jose M Such and Michael Rovatsos. 2016. Privacy policy negotiation in social media. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* 11, 1 (2016), 1–29.
- [37] Onuralp Ulusoy and Pinar Yolum. 2021. PANOLA: A Personal Assistant for Supporting Users in Preserving Privacy. *ACM Transactions on Internet Technology* 22, 1, Article 27 (sep 2021), 32 pages.