

What is in your cookie box? Explaining ingredients of web cookies with knowledge graphs

Geni Bushati ^{a,c,*}, Sven Carsten Rasmusen ^a, Anelia Kurteva ^{a,**}, Anurag Vats ^b, Petraq Nako ^c and Anna Fensel ^{a,d,e}

^a *Semantic Technology Institute (STI) Innsbruck, Department of Computer Science, University of Innsbruck, Innsbruck, Austria*

E-mails: Geni.Bushati@student.uibk.ac.at, sven.rasmusen@sti2.at, a.kurteva@tudelft.nl

^b *Distributed and Parallel Systems Group (DPS), Department of Computer Science, University of Innsbruck, Innsbruck, Austria*

E-mail: anurag.vats@uibk.ac.at

^c *University of Innsbruck, Innsbruck, Austria*

E-mail: Petraq.Nako@student.uibk.ac.at

^d *Wageningen Data Competence Center, Wageningen University & Research, Wageningen, The Netherlands*

^e *Consumption and Healthy Lifestyles Chair Group, Wageningen University & Research, Wageningen, The Netherlands*

E-mail: anna.fensel@wur.nl

Editors: Bo Fu, California State University Long Beach, USA; Patrick Lambrix, Linköping University and University of Gävle, Sweden; Catia Pesquita, LASIGE, Faculdade de Ciências, Universidade de Lisboa, Portugal

Solicited review: Three anonymous reviewers

Abstract. The General Data Protection Regulation (GDPR) has imposed strict requirements for data sharing, one of which is informed consent. A common way to request consent online is via cookies. However, commonly, users accept online cookies being unaware of the meaning of the given consent and the following implications. Once consent is given, the cookie “disappears”, and one forgets that consent was given in the first place. Retrieving cookies and consent logs becomes challenging, as most information is stored in the specific Internet browser’s logs. To make users aware of the data sharing implied by cookie consent and to support transparency and traceability within systems, we present a knowledge graph (KG) based tool for personalised cookie consent information visualisation. The KG is based on the OntoCookie ontology, which models cookies in a machine-readable format and supports data interpretability across domains. Evaluation results confirm that the users’ comprehension of the data shared through cookies is vague and insufficient. Furthermore, our work has resulted in an increase of 47.5% in the users’ willingness to be cautious when viewing cookie banners before giving consent. These and other evaluation results confirm that our cookie data visualisation approach and tool help to increase users’ awareness of cookies and data sharing.

Keywords: Cookies, consent, GDPR, ontology, knowledge graph, data sharing, comprehension

*Corresponding author. E-mail: Geni.Bushati@student.uibk.ac.at.

**The author is currently affiliated with Delft University of Technology, Delft, The Netherlands.

1. Introduction

Cookies have emerged as one of the most convenient and common mediums to request consent for data sharing online [55]. The rising digitisation of services in the e-commerce, healthcare, finance and social media domains have also turned cookies into a valuable source of personal data such as IP address and browsing behaviour. Cookies are a promise for better user experience online as the data that they collect is often used for user profiling to create personalised browsing experiences and e-commerce recommendations. However, this is often done at the cost of an individual's privacy [27,41]. Identifying who benefits more from cookies, individuals sharing their data or companies that use it, has also become a challenge. The European Union's (EU) GDPR [17], in effect since May 2018, has highlighted the importance of consent and has set it as one of its legal basis for personal data processing. Consent must be freely given, specific, informed, unambiguous and one should be able to revoke it with the same ease it was given (Art. 4 (11)). As shown in [4,32], however, requesting and receiving informed consent does not equal individuals being truly aware of the implied data sharing. Requesting informed consent prior to any data processing can cause consent fatigue [57] and information overload [26], which often lead to blindly given consent [3]. Several factors, namely the lack of knowledge of what cookies are and of their functions [36], the lack of control that users have over the data that is collected and shared through cookies and the lack of feedback provided by browsers' cookie management facilities [24] contribute to this.

Despite GDPR's requirements for informed consent, many cookies are still not compliant with regard to the information (or lack thereof) that they present to individuals and how they are imposed [35]. Further, once consent is granted, the cookie dialogue disappears (i.e., it is no longer visible while browsing) and data sharing begins on the back-end. This also poses a challenge for revoking consent (Art 7(3)) as there is no cookie dialogue via which individuals can exercise their right to personal data erasure (Art. 17(2)) [50]. Many online service providers can benefit from the disappearing cookie dialogues as without a reminder cookies are likely to be forgotten by individuals. Retrieving cookie logs, which store specific data about the consent, data processing and the cookie's duration, can be a complex task for individuals with no prior privacy or technical experience (i.e., cookies logs are described using privacy and security terminology). Several solutions, in the form of browser extensions such as the Cookie Editor¹ (utilised in this paper) ease the retrieval of existing cookies by allowing users to directly export cookies instead of examining the browser's logs to locate them. Such tools focus mainly on simplifying the cookie retrieval process, leaving behind the meaning (semantics) of the stored cookies and are browser-dependent (e.g., built for Google Chrome²), which limits their wider utilisation.

There is a need for greater clarity, transparency and awareness about cookies and the data sharing that happens behind the scenes. Both the design of cookie dialogues and the triggered data sharing need to comply with GDPR (when the individual is an EU citizen). Achieving this, however, is challenging as user interface (UI) and user experience (UX) designers might not have the same legal expertise as a data protection practitioner. A single consistent schema outlining all the information that cookies need to present to individuals based on the GDPR can be used to harmonise the domain experts' work. Further, such schema can help to ensure legal compliance and can bring more transparency into cookie-based data sharing. This can be achieved with Semantic Web technologies such as ontologies, which represent a domain in a machine-interpretable format and can be used as a schema for knowledge graphs (KGs) to further interlink multiple domains [20]. KGs also support data interoperability, transparency and traceability [13,21] and are extendable by design, which makes them suitable for use in different ecosystems and across multiple use cases. In the security and privacy domains, KGs have been successfully utilised for privacy-enabled penalisation on the web [25], intelligent decision-making, fraud detection, prediction and tracing of cyber attacks (see [28,42,44,60]). Other domains such as manufacturing (e.g., [9]) and logarithmic law (e.g., [10,54]) have also significantly benefited from utilising KGs to bridge knowledge silos, semantically enrich data, highlight data dependencies and discover insights and new knowledge. Multiple ontologies for data sharing such as Consent and Data Management Model (CDMM) [19], Data Protection Vocabulary (DPV)³ [40] and GConsent [39] have been built and are widely utilised [33]. However, there is currently no ontology for web cookies in the context of GDPR

¹<https://cookie-editor.cgagnier.ca>

²<https://www.google.com/chrome/>

³<https://w3id.org/dpv>

or beyond it, which can guide the standardisation of cookie consent dialogues' design. The existing ontologies focus primarily on the concepts of consent, contracts, or data sharing and lack the semantic representation of web cookies.

Motivated by this and by building upon the findings in [4,46] that highlight the need for greater online data sharing transparency and interpretability, we present the OntoCookie^{4,5} ontology for machine-readable and standardised cookie representation and a KG-base tool⁶ for personal cookie visualisations built with it. The main goal of our work is to bring more transparency and awareness regarding cookies and to ease individuals' comprehension of cookie-based data sharing. Further, we believe that our ontology sets the groundwork for the future synchronisation of the design, legal and technology domains in the case of data sharing. The main research question that we answer is: "Can a KG-based visualisation of cookie statistics help to ease one's comprehension of cookie data sharing?". In the context of this paper, the ease of cookie comprehension refers to the ability of users to understand what exactly a cookie is (i.e., its source, duration and type). To summarise, we make the following contributions:

- A novel, open-access documented OWL ontology for cookies (referred to as OntoCookie⁴).
- Cookie insights derived from the analysis of the KG with cookies of 40 users.
- A novel, open-access tool for personalised cookie visualisations that has been evaluated with 40 users.*

The rest of the paper is structured as follows. Section 2 presents an overview of related work relevant to our study. Section 3 outlines our approach and the followed methodology. The implementation of our work is presented in Section 4, while its evaluation and its results are presented in Sections 5.1 and 5.2 respectively. A summary and discussion on results is presented in Section 5.3. Conclusion and future work are presented in Section 6.

2. Related work

This section presents related work on cookies as an online medium for consent from the privacy, visualisation and Semantic Web fields that helped to motivate our work.

2.1. Cookies and privacy

For many years, cookies have been viewed as a privacy-preserving mechanism [30]. However, the enforcement of the GDPR and its requirements for the lawful processing of personal data have highlighted the numerous privacy risks associated with them. According to Article 4(1) and Recital 30 of the GDPR, cookies and specifically cookie identifiers are viewed as personal data, which needs to be handled in compliance with the law. This also includes the need for informed consent request for each cookie. Santos et al. [51] present an in-depth analysis of how data-sharing information is presented via cookie consent dialogues. Following the legal requirements of the ePrivacy Directive (ePD) [18] and the GDPR, around 400 cookie banners presented on the most popular English-speaking websites were manually annotated. 89% of the cookie banners violated the applicable laws. More specifically, 61% of the banners violated the purpose specificity requirement by mentioning vague purposes, including "user experience enhancement" while further, 30% of banners used positive framing, breaching the freely given and informed consent requirements. In a similar study, Soe et al. [52] analysed 300 data collection consent notices from news outlets, which were built to ensure GDPR compliance. The analysis uncovered the use of a variety of dark patterns (i.e., deceptive design practices aimed at manipulating users' actions) [22,34].

Sanchez-Rola et al. [49] explore users' perception and reaction to cookie dialogues and conclude that users view cookie dialogues as an annoyance during their browsing time rather than an informative source. Although the users claimed to have privacy concerns regarding cookies and how they collect data, the study showed that the cookie disclaimers did not play a significant role in the users' decision to continue navigating the website. Greater

⁴<https://genibushati.github.io/OntoCookie.io/>

⁵<https://github.com/STIIInnsbruck/OntoCookie/>

⁶<https://svencarstenrasmusen.github.io/cookie-consent-visualisation-tool/build/web/index.html#/>

*Currently, the online tool is functional only when clicking "No. I disagree" to the consent banner.

importance was given to factors such as the reputation of the website, which can also affect the users' trust in its services [12,23].

In a similar study, Bechmann [3] shows that there exists a non-informed consent culture among social media platform users and that although none of the participants of the study had read the privacy policies, all have given consent. Joergensen et al. [29] further confirm that users rarely read the presented data-sharing terms and conditions before granting consent. Furthermore, statistics from countries within and outside the EU show that most users of social networking sites do not read the privacy policies of the sites or the third-party applications that use their data [3]. The studies confirm that for users, giving consent (in any form such as cookies) and being aware of what the action implies are often mutually exclusive [3,29,49].

2.2. *Cookie visualisation*

The lack of transparency about what accepting a cookie implies and the lack of accessible information about it further contribute to blindly given consent online. Ware [58], Rossi et al. [47] and Drozd et al. [16] highlight the importance of visualisation as a way to support the comprehension of the information that is being communicated to the end user. According to [58], the highest bandwidth channel of communication between humans and machines is provided by visual displays. The amount of information that can be transmitted makes data visualisation a highly appropriate method to communicate information to users.

Rossi et al. [47] emphasize the fact that the use of visualisation is explicitly suggested by the European Union (EU) in legislations such as the GDPR (Rec. 58, Art. 12(7)) as a way to improve comprehension of the information provided to data subjects. One can acknowledge that visual elements and visualisations in general play a crucial role in obtaining informed consent. In recent years there has been a rise in the attempts to build applications that provide more transparency regarding personal data processing through applying different visualisation approaches.

Steichen et al. [53] go deeper into the topic of information visualisation by taking into consideration the role of the individual cognitive style of the users in their ability to perceive the information being communicated in a visual form. Results show that the individual cognitive style plays a significant role in tasks related to information visualisation in general. Findings of the presented work also provide motivation for the development of personalised information visualisation systems based on the cognitive style of the individual users.

In this context, Drozd et al. [16] present the CoRe [15] and the Consent reqUest useR intErface (CURE) user interfaces (UIs), which have the main goal of easing the process of granting consent and providing more transparency into data sharing. The evaluation of the two UIs showed that, indeed, visualisations helped raise awareness of what consent is. However, issues such as information overload due to design complexity were still present. Similar solutions for consent visualisation include the work in [1] and [45], which focuses on raising data-sharing awareness with visualisations. All these studies show that there is a prominent need for consent solutions that support higher levels of transparency, focus on the needs of the users and on raising awareness regarding data sharing.

2.3. *Cookies and the Semantic Web*

One of the earliest and few studies on cookies through the lens of semantics is presented by Cox et al. [11]. The authors explore the application of the Semantic Web in the privacy field and propose an approach for enriching cookies with Resource Description Framework (RDF) fragments. The main goal of the created semantic cookies is to ease access to web services and give users full control over their data online while widening their participation in the Semantic Web. The study has shown promising results and highlights the benefits of machine-readable cookies for persistence stores to simplify access to services. Cox et al. are one of the first to discuss the use of an ontology as a tool that can align the various representations of cookies online, which can also support legal compliance. However, we were not able to identify any such existing publicly available ontology. Through the years most of the work has focused predominantly on consent. A systematic analysis of semantic models for consent and semantic-based visualisations tools that supports users' comprehension of consent is presented in [33].

A more recent work that addressed cookies and online data-sharing privacy policies is presented in [2]. Audich et al. [2] propose an ontology for privacy policies, which also includes the concept of a cookie and combines it with natural language processing. The main goal of the approach is to improve the readability of online policies

by identifying the key information in a policy for individuals to focus on. Cookies and instances such as do-not-track and web beacons (types of cookies) have been semantically represented as keywords that can be found within policy documents or cookie policies. The results of the study have proven the benefits of utilising an ontology (i.e., helping to align and simplify the complex and diverse legalese that is used) for text mining of privacy policies [2]. The study focused on privacy policies in general and legal terminology used by the Federal Trade Commission.⁷ Specific legislation such as the GDPR and visualisation as a tool to increase transparency and ease comprehension have not been explored.

Motivated by the lack of consent interoperability and transparency regarding data sharing online, Bless et al. [4] utilise both semantics (i.e., ontologies and KGs) and visualisations as key tools to support individuals' comprehension of consent for data sharing. In comparison to the work in [15] and [16], the authors focus on visualising data-sharing flows after consent is given. The developed visualisation is based on a KG, which stores informed consent information in a GDPR-compliant manner and has helped to raise individuals' awareness of data sharing significantly. Further, as shown in [10], the later version of the consent KG has also been successfully utilised for performing GDPR compliance verification and in supporting humans and machines in making sense of consent [32].

The use of ontologies for consent and GDPR compliance is also prominent in the work of Kirrane et al. [31], that shows success in utilisation of semantics to build more accurate models to detect security issues. Moreover, the meaningful interpretation of personal data that is exchanged between users and other entities on the web can be used to empower users to have better control over these interactions and therefore improve the way they manage their online privacy. The semantic approach can also bring advantages to companies through automation, which is enabled by the semantic machine-readable and machine-processable representation of data-related privacy policies. The main trends for utilisation of KGs in the security and privacy domains are further discussed in [10,31]. The benefits of semantics in the legal domain, especially for improving consent interoperability, are also discussed in [8,20,33,59].

Rasmusen et al. [46] present a KG-based interface that visualises consent request and utilises gamification to raise user engagement in data sharing. The main goal of the approach is to improve individual's awareness of consent and the implications that follow in the context of automotive data. The UI presented follows an ontology that models GDPR knowledge about consent. Results from the user study conducted with participants that interact with the tool show that the UI helped raise the individuals awareness and willingness to consent.

2.4. Summary

The GDPR has set out specific requirements for requesting consent in an informed way through any medium including cookie dialogues. However, research has shown that there is a lack of standardisation with regard to the design of cookie dialogues and the information presented on them. There is currently a misalignment between law, technology and design when it comes to cookies and the underlying personal data that is collected and shared via them. The proposed work in [11] and [2] has called attention to the benefits of semantics in the privacy field concerning cookies. To our knowledge, there is currently no publicly available vocabulary or ontology that can align the knowledge spread across these domains in the context of GDPR. Cookies have become the go-to tool for many service providers when it comes to personal data collection online. Although this has raised privacy concerns due to the lack of transparency of cookie-based data collection and sharing, there is a lack of user-centered tools that support the comprehension of what cookies are and the implications of giving consent for them. To summarise, based on our research of related work, two main challenges have become evident – the lack of shared vocabulary of the cookie domain that can support knowledge exchange and data interoperability for legal compliance and the lack of support for users in making sense of cookie-based data sharing.

3. Selected approach and methodology

We approach the issue of web cookie comprehension and cookie data sharing from both human and machine perspectives. However, both sides have different comprehension needs that need to be addressed. On the human

⁷<https://www.ftc.gov>




Fig. 1. Methodology overview.

side, we focus on utilising data visualisations in graphical and tabular forms. Our cookie visualisation tool provides individuals with an interface that takes as input cookie logs and displays personalised statistics that are aimed at providing more transparency into cookie-based data sharing. Consequently, this can help raise individuals' awareness regarding the implications of granting consent for cookies. The OntoCookie^{4,5} ontology and the KG built with it are used to represent the cookie data in a meaningful machine-readable and interoperable way. Our approach is motivated by the increase of cookie and consent requests online after the acceptance of the GDPR and tries to bridge the gap between the Semantic Web, privacy and legal domains.

The methodology (Fig. 1) followed for the development of our cookie visualisation tool is inspired by the design thinking process [6], which is a solution-based approach to solving problems by considering human needs. The development process consisted of the following stages: emphasise, define, ideate, prototype and test. The first stage was to understand the problem of cookies and consent comprehension. This included research on the privacy domain and, more specifically, on cookies and how data and consent are handled by browsers. Existing work on cookies (with and without the use of semantics) was also considered (see Section 2). During the second stage, the main research problem was defined and system requirements were derived. The third stage focused on analysing the requirements and generating ideas for the design of the tool. The fourth stage focused on prototyping the solution. This was done in several stages as well. We started with (i) building the OntoCookie ontology, (ii) building a prototype UI for cookie import, (iii) implementing functionalities such as cookie annotation, (iv) building the cookie KG and finally (v) visualising different cookie statistics on the UI. The fifth stage consisted of the usability and design evaluation of the tool with users, analysis of the results and the comparison to existing cookie solutions. Our cookie visualisation tool was built with the Flutter⁸ toolkit for front-end development, NodeJS⁹ on the back-end environment, Protégé¹⁰ and GraphDB¹¹ for building and storing the OntoCookie ontology and KG. The Cookie Editor¹ browser plug-in (available for Google Chrome, Firefox, Opera and Microsoft Edge) was used to allow users to export their cookies for each website separately.

4. Implementation

This section presents details regarding the implementation of the proposed KG-based tool for cookie visualisations. Section 4.1 presents an overview of the OntoCookie ontology for cookies, which has been built and utilised during the study. Section 4.2 presents the two possible action flows of using our tool, while Section 4.3 presents the implementation details of the visualisation.

4.1. OntoCookie: A domain ontology for cookies

The OntoCookie^{4,5} ontology (Fig. 2) is a formal representation of the cookie domain in the context of GDPR. The ontology was built as a response to the lack of openly available semantic models for cookies and the need for cookie consent compliance (from a design and implementation perspective). By following a top-down ontology engineering approach (see [37]), the main classes, sub-classes the relationships between them and their data properties

⁸<https://flutter.dev>

⁹<https://nodejs.org/en/>

¹⁰<https://protege.stanford.edu>

¹¹<https://www.ontotext.com/products/graphdb/>

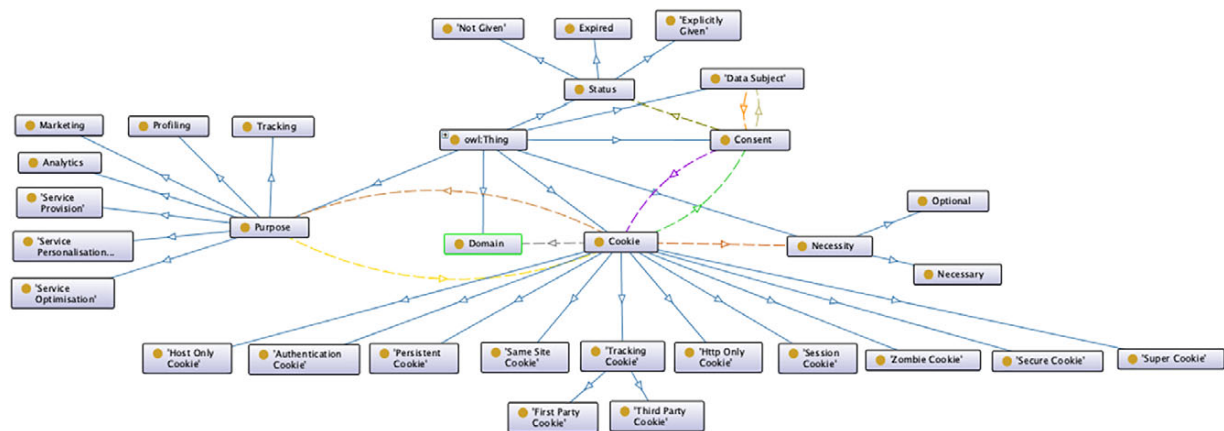


Fig. 2. The OntoCookie ontology.

were defined. When defining the subclasses, an “isA” constraint was followed (e.g., *SessionCookie* isA *Cookie*). OntoCookie was built with Protégé¹² and currently comprises of 229 axioms, 32 classes, 10 object properties and 10 data properties. The latest version (version 1.2) of the ontology is publicly available in GitHub⁵ and has been documented⁴. This version of the ontology was evaluated with the HermiT¹³ reasoner for inconsistencies and with the OOPS! [43] ontology pitfall scanner.

The class *OntoCookie:Cookie* represents several types of cookies that are widely used (e.g., session, host only, HTTP only, persistent, authentication, tracking). Definitions for each cookie type have been provided as well by reusing *dc:description* from the Dublin Core¹⁴ vocabulary. To model metadata such as the *startDate* and the *endDate* (in an ISO 8601¹⁵ format) of a cookie, the data property *schema:Date* was reused. Further, the data property *schema:Duration* can be used to represent the duration of a specific cookie. Cookies can also be related to the specific web domain (*OntoCookie:Domain*) for which they are valid. If a domain is not specified, then the hostname of the originating server is used as the default value. The necessity of a cookie can be represented as well. *OntoCookie:Necessary* cookies are essential for a service to function, while *OntoCookie:Optional* can be used to collect additional data for various purposes such as *OntoCookie:Analytics*, *OntoCookie:Marketing*, *OntoCookie:Profiling*. Linking a cookie to its purpose can further support GDPR compliance verification as consent is represented by *gconsent:Consent* (requested through any medium) should be informed and should have a specific purpose. A cookie and the consent for it are given by a specific *OntoCookie:DataSubject*. While using the cookie visualisation tool, each user is asked to generate a unique hash (modeled by the data property *OntoCookie:hashed_id*), which is used later for retrieving the specific cookies from the generated KG.

4.2. The action flow

In order to adhere to GDPR regulations, the users are asked for their informed consent (i.e., users are explicitly asked whether they want their cookie data saved in the KG via a consent dialog). If consent is given, the action flow consists of 11 steps (numbered from 1 to 11 in Fig. 3). In step 1, the users are provided with details about the application. In step 2, the users can import the collected cookie data in JSON format into a designated text field for this purpose. Next, the users have the option to visualise their data. However, before continuing to the next step they are asked for consent via a consent dialog.

In the case of consent, the action flow continues to step 3, where the JSON file with cookie data is directed to the NodeJS middleware. A SPARQL INSERT query is constructed by the middleware in step 4, where it is executed

¹²<https://protege.stanford.edu/>

¹³<http://www.hermit-reasoner.com>

¹⁴<https://www.dublincore.org>

¹⁵<https://www.iso.org/iso-8601-date-and-time-format.html>

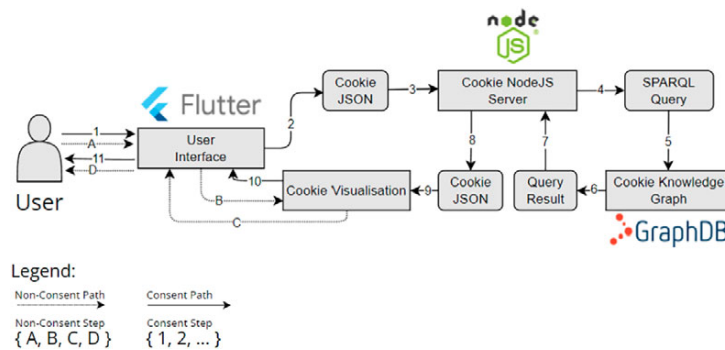


Fig. 3. The action flow of the application.

in the KG during step 5. Upon completion of the query, the middleware executes a `SELECT` query in steps 6 and 7, where it fetches all the uploaded cookie data by the current user. In step 8, the queried data is stored in a JSON format. In step 9, the data is processed and consequently visualised in step 10. In step 11, the users can view their cookies data in a human-readable format.

In case the users do not consent to have the cookies stored in the KG, the data will not be annotated to the KG. Users can import the cookies collected into the designated text field created for this purpose. After deciding not to consent (i.e., users have decided not to save the cookie data consumed in the KG), the data will be locally processed and accordingly visualised. In this case, cookie data imported through the Cookie Editor extension will be deleted once the application window is closed. The non-consent action flow follows steps A, B, C and D in succession (Fig. 3).

4.3. UI and the connection to the back-end

The UI is organised in two parts. The first part is a general guide of six steps on how to use the visualisation tool (Fig. 4). Step 1 contains a link to the extension we use to export cookies in JSON format. Step 2 asks the user to enter their randomly generated ID, which is created at the start of our evaluation process. Steps 3 and 4 explain how the user should use the Cookie Editor browser extension to import their cookies into the visualisation tool. For each website, a separate import has to be done, as the browser extension loses the cookies of a website once the user navigates to a different website domain (e.g., navigating from “Wikipedia.org” to “Euronews.com”). Once the users click on the *Visualise* button, they are asked if they would like to consent to store their cookies on the KG for 10 days, as explained in Section 4.2.

Consequently, the second part of the UI displays all the cookie data, except the stored value, retrieved from the browser extension (Fig. 5). Here, the information is divided into four segments. Segment 1 lists all cookies with their domain, name, type and duration. Segment 2 contains a bar chart grouping the amount of cookies based on their duration. A pie chart containing the distribution of the cookies among all visited websites is visualised in segment 3. Charts were created with the help of the `charts_flutter`¹⁶ library. Segment 4 contains a button that gives the users the opportunity to withdraw consent and erase data from the KG if they agreed to share it previously. Segment 5 illustrates which websites stand out for storing cookies (i.e., longest cookie, shortest cookie, the total amount of cookies and the average duration of all cookies combined).

To build the back-end, we used NodeJS and also Express for the routing. This has made the creation of our application programming interface easy to use. As mentioned in previous sections, we have created a KG in order to save the information on cookies and their relations with each other. Our KG is contained in GraphDB, a graph database for KGs in RDF. The back-end is connected to GraphDB using the `sparql-client` library. In this way, we perform `SELECT`, `INSERT` and `DELETE` queries against a SPARQL endpoint via HTTP. All the source code can be found at our GitHub⁵ including the OntoCookie ontology and a link to try out the cookie visualisation tool.

¹⁶https://pub.dev/packages/charts_flutter

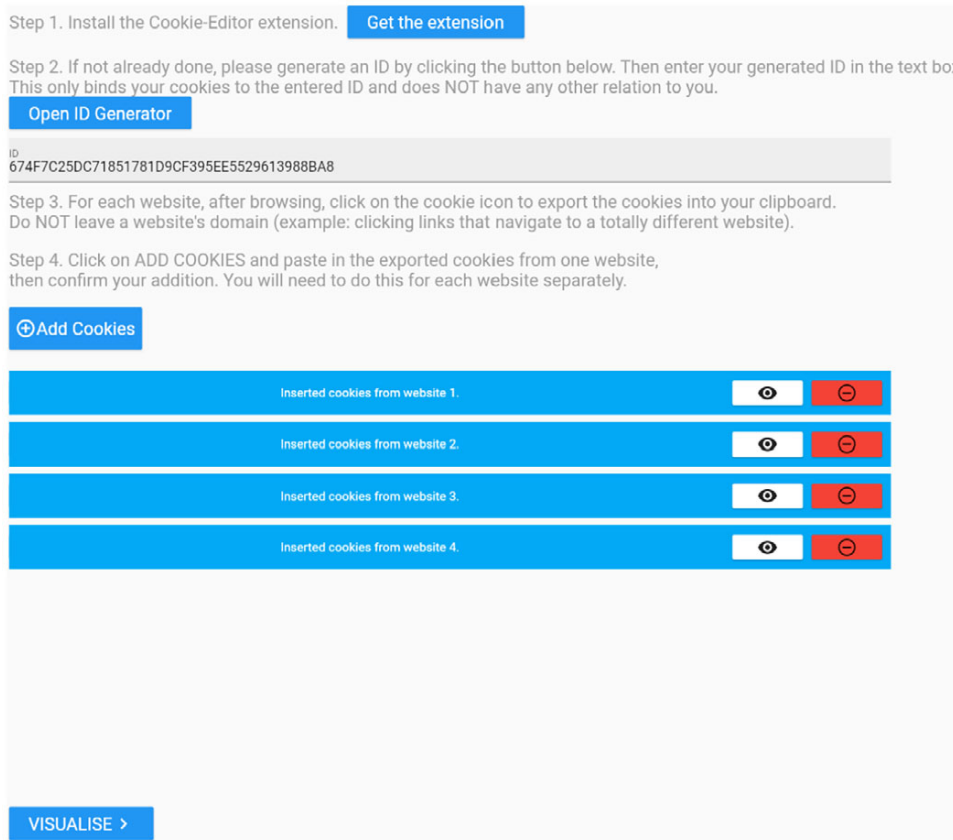


Fig. 4. Main input page of the cookie visualisation tool.

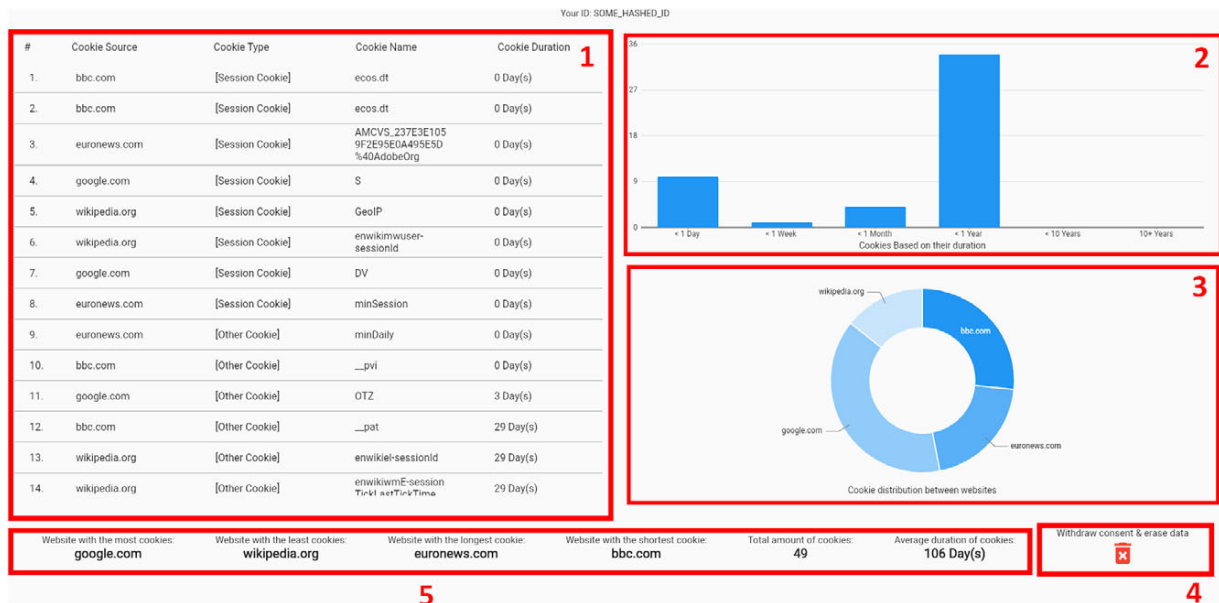


Fig. 5. Overview of the cookie visualisation statistics.




Fig. 6. Evaluation flow.

In order to get the time of execution of these SPARQL queries,¹⁷ we ran them on GraphDB. The `SELECT` and `DELETE` queries were executed with 70 cookie instances (i.e., the average amount of cookies collected by a single user while performing the evaluation). The queries were performed on a regular personal laptop and the GraphDB in this case was hosted on a server. Execution time for `SELECT` and `DELETE` queries was 0.1 s for each query.

5. Evaluation

This section presents details about the evaluation of the presented cookie visualisation tool, namely the evaluation set up (Section 5.1), evaluation results (Section 5.2) and a summary of these results (Section 5.3).

5.1. Evaluation set up

To evaluate our solution, its usability and design, three questionnaires (on demographics, expectation, and realisation) were created using Google Forms. The evaluation was done in seven stages (Fig. 6). First, the participants were asked to generate a unique ID using the SHA1¹⁸ online hash generator and then to complete a demographics survey. Next, in stage 3, they were presented with an introductory video¹⁹ that contains general information on cookies (what cookies are, different cookie types, etc.). The goal of this was to familiarise the end-user with the topic. Following the video, in stage 4, the participants were asked to install the Cookie Editor extension and to browse four websites for a time span of two minutes. The extension is available for the Google Chrome, Firefox and Microsoft Edge browsers and provides an export button that allows the users to export their cookie data into their clipboard. For the work with the cookies collection, we have selected four highly used websites (“Google.com”,²⁰ “Wikipedia.org”,²¹ “BBC.com”,²² and “Euronews.com”²³) that do not require users to register to access information. In this way, the cookies which we collect do not have sensitive information such as usernames and passwords. During stage 5, the participants were presented with a pre-use (i.e., before using the tool) expectation survey, which contains questions to evaluate their general knowledge of cookies and the expectation of what data cookies can collect. Having completed that, in stage 6, the participants were asked to export their cookies with the Cookie Editor and to import them into the cookie visualisation tool and visualise the data. To measure in a quantified manner whether participants’ comprehension of cookies has changed after using the presented tool, all participants were asked to fill in a post-use realisation survey. The analysis of the results is presented in the next sections.

5.2. Evaluation results

For the evaluation, 40 participants (25 male and 15 female) took part in the survey. The age of the participants varied between 18–35 years old where 92.5% were within the range of 18–30 years old and 7.5% were within the range of 30–35 years old. The participants were selected from different backgrounds (computer science students,

¹⁷https://github.com/STIIInnsbruck/OntoCookie/tree/main/tool/cookie_server/controllers

¹⁸<https://passwordsgenerator.net/sha1-hash-generator/>

¹⁹<https://www.youtube.com/watch?v=KKZIEaAWAao>

²⁰<https://www.google.com/>

²¹<https://www.wikipedia.org/>

²²<https://www.bbc.com>

²³<https://www.euronews.com>

non-computer science students, researchers, computer-science experts, non-computer science experts) and were based in different countries in Europe (namely, Austria, the Netherlands, Bulgaria and Albania). They were recruited via university network and personal connections. Out of the 40 participants, 20% acknowledged that their highest level of education completed was a high school degree. For 57.5%, the highest level of education obtained was a bachelor's degree and for 22.5%, the highest level of education obtained was a master's degree. 30% of the participants declared a very high-level Internet surfing competency, 47.5% a high level of competency, 20% declared an average Internet surfing level competency and 2.5% declared a low level of Internet surfing competency. 65% of our participants spend more than 4 hours per day on the Internet, 25% spend 3–4 hours per day and 10% spend 1–2 hours per day on the Internet.

5.2.1. Expectation vs. realisation

In order to measure the level of comprehension of the users in regards to the cookies collected during the browsing time of the four websites, we at first asked them about their expectation (i.e., how the users expected the results to be before using the application) and compared them with the personalised data (i.e., the factual results) which were visualised by the application. For this purpose, questions related to the amount, duration and source of the cookies collected were asked.

More specifically, to the question: “Which of the websites do you think has the highest amount of cookies?”, 80% of our participants expected it to be “Google.com”, 2.5% expected it to be “Wikipedia.org”, 5% expected “Euronews.com” and 12.5% expected “BBC.com” to have the highest of cookies (Fig. 7(a)). In contrast to the users' expectations, the results showed that for 60% of the participants, the highest amount of cookies consumed was generated from “Euronews.com”, for only 30% the highest amount of cookies collected was from “Google.com” and for 10%, the highest amount of cookies collected was from “BBC.com” (Fig. 7(b)).

To the question: “Which of the websites do you think has the least amount of cookies?”, 55% of the participants expected it to be the website “Wikipedia.org”, 27.5% expected it to be “Euronews.com”, 10% expected it to be “Google.com” and 7.5% expected to be “BBC.com” (Fig. 7(c)). Results showed that in 77.5% of the cases, “Wikipedia.org” had the least amount of cookies, in 10% of the cases “BBC.com” had the least amount of cookies collected, followed by “Google.com” and “Euronews.com” with 7.5% and 5% respectively (Fig. 7(d)).

When asked: “Which of the websites do you think has the longest lasting cookie?”, 80% of the participants expected it to be “Google.com”, 12.5% expected it to be “Wikipedia.org” and 7.5% expected that the longest lasting cookie originated from “BBC.com” (Fig. 8(a)). Meanwhile, the results obtained show that in 52.5% of the cases, “Google.com” had the longest lasting cookies, “BBC.com” had the longest lasting cookie in 27.5% of the cases and on 20% of the cases the longest lasting cookie belonged to “Euronews.com” (Fig. 8(b)).

To the question: “Which website do you think has the shortest lasting cookie?”, 47.5% of the participants expected “Wikipedia.org” to have the shortest lasting cookie, 35% answered “Euronews.com”, 12.5% answered “BBC.com” and 5% expected the shortest lasting cookie to belong to “Google.com” (Fig. 8(c)). On the contrary, the realisation results showed that in 67.5% of the cases, “BBC.com” had the shortest lasting cookie, in 17.5% of the cases “Euronews.com” had the shortest lasting cookie, in 10% of the cases “Wikipedia.org” had the shortest lasting cookie while on 5% of the cases the shortest lasting cookie belonged to “Google.com” (Fig. 8(d)).

The claim that the users' knowledge on cookie data is vague and insufficient was further strengthened by the significant differences detected between expectation and realisation, with respect to the total amount of cookies

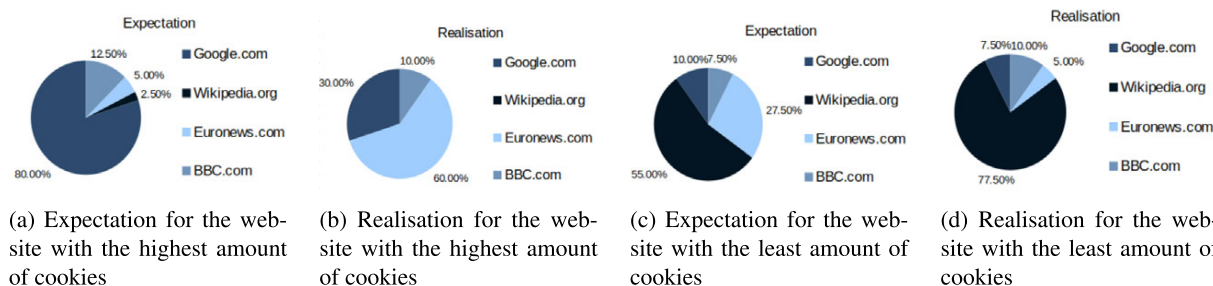


Fig. 7. Survey results for the amount of cookies collected by each website.




Fig. 8. Survey results for the duration of cookies collected by each website.

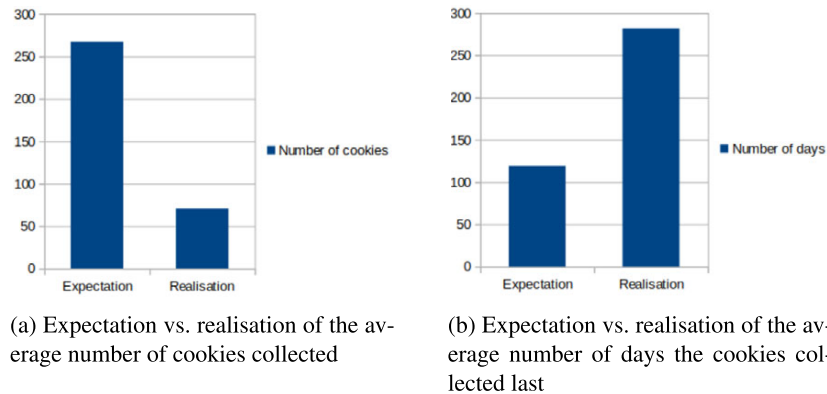


Fig. 9. Comparison of expectation vs. realisation averages.

collected and their duration. More precisely, on average, the participants expected the total number of cookies collected during the two minutes of website browsing to be 267.4. Results from the realisation survey showed that, on average, a total amount of 70.8 cookies were collected during their surfing time, approximately 73% less than the users' expectation (Fig. 9(a)). Regarding the duration of cookies collected, when asked: "How many days on average do you think cookies last?", the response mean was 119.2 days. Results obtained from the realisation survey show that on average, the cookies collected during the session lasted 281.8 days, approximately 137% more than the expectation (Fig. 9(b)).

The question: "How carefully do you read the cookie notification banner before proceeding to give consent or not?" was numerically encoded on a scale from 1 ("Not carefully at all") to 5 ("Very carefully") and was asked to the participants before using the application. 82.5% said that they do not read the banner carefully at all or not carefully, while 17.5% were neutral, read the banner carefully, or very carefully (Fig. 10(a)). After the participants used the application, we asked the question: "How carefully will you be reading the cookie notification banner before agreeing to give consent or not?". Participants responded that they were willing to be more careful when reading the cookie notification banner before agreeing to cookies, showing a significant increase in awareness related to the process of web cookie agreement. Specifically, 65% of the participants were neutral, willing to be careful, or willing to read very carefully the cookie notification banner and 35% of the participants confessed that they would continue not to be careful or not carefully at all when agreeing to the cookie notification banner (Fig. 10(b)).

5.2.2. Further survey results

Furthermore, participants answered a set of questions related to their general feeling about cookie data privacy after using the application and also how they will approach Internet cookies in the future. The participants had the perception that cookies were intrusive to their online privacy. Specifically, to the question: "Do you feel as if the website knows more than you expect", 62.5% of the participants answered "Yes" while 37.5% answered "No". The possible answer to the question: "How do you feel about your privacy when browsing the Internet in regards to the safety of your data, after being presented with information on the cookies you consumed?" was numerically encoded on a scale from 1 ("Not safe at all") to 5 ("Very safe"). 57.5% of the participants replied either 1 or 2,




Fig. 10. Comparison of results regarding the carefulness with which the participants read and will read the cookie notification banner before and after using the application.

meaning that they did not feel safe about their data privacy, 37.5% were neutral, while 5% felt safe in regards to their data privacy on the Internet. To the question: "Do you feel it is fine for websites to collect the given amount of cookies?", 82.5% of the participants answered "No" and that they "Wished for fewer cookies to be collected", 15% answered "Yes" and that "Things may continue unchanged", and 2.5% answered "No" and that they "Wished for more cookies to be collected".

Results showed that participants would embrace an overview tool to manage their cookies. Precisely, the question "Would you feel more confident surfing the Internet if you were given an overview tool to manage your cookies?". 72.5% answered "Yes", 25% were neutral and 2.5% answered "No". Further, we asked the participants: "Would you feel more knowledgeable about cookies and your browsing privacy if you were given an overview tool to manage your cookies?". Results show that 95% of the participants would feel more knowledgeable about cookies and browsing privacy if they were given an overview tool to manage them. 2.5% were neutral while 2.5% responded "No".

5.3. Summary of the results

Evaluation results confirm that users (even proficient web surfers) lack detailed knowledge about cookies and the consequences of granting consent for them. For example, the duration of the cookies being stored, the amount of cookies collected during the browsing time and practices of different websites with regard to the cookies they use, commonly do not match the users' expectations. The results also showed that the cookie visualisation tool presented helped to improve users' comprehension of cookies and has raised awareness regarding data sharing on the web. More specifically, after being presented with the application, an increase of 47.5% in the users' willingness to be more cautious when reading the cookie consent banner before giving consent was noticed. The outcome of the evaluation also confirms that users are ready to embrace an overview tool that helps them manage their cookies. 72.5% of the participants agreed that they would feel more confident about their privacy on the web if they were given such overview tool and 95% of the users admitted that they would feel more knowledgeable about cookies if an overview tool to manage cookies was at their disposal. In addition, we believe that this work helps breach the gap between the Semantic Web and the security and privacy domains.

Table 1 describes how our work compares with the existing work in this field in several aspects. It contains information on results, consent request medium, use of semantics, whether the work focuses on before or after consent is given, whether it includes usage of cookies, and lastly, limitations. Results from the table confirm that there is currently a lack of semantic approaches that describe online cookies in order to enhance the users' understanding of the cookie data they consume on a personal level while surfing the Internet.

Table 1
Comparison of existing solutions for online consent request

Study	Result	Consent request medium	Use of semantics	Pre- or post-consent stage	Focus on cookies	Limitations
Cox et al. [11]	Enriching cookies with RDF.	Cookies	Yes	Post-consent	Yes	The security and privacy of exchanging RDF cookies have not been addressed.
Drozd and Kirrane [15]	UI for personalised consent requests.	Consent form	No	Pre-consent	No	Complex visualisation causing information overload.
Bless et al. [4]	UI for post-consent data flow visualisation.	None	Yes	Post-consent	No	The solution focuses only on visualising consent data flows. Exercising users' rights is performed via a different tool (see Rasmusen et al. [46]).
Audich et al. [2]	A tool for minimising privacy policies.	None	Yes	Pre-consent	No	The considered privacy policies do not consider the GDPR legislation.
Rasmusen et al. [46]	UI for consent request that utilises gamification.	Consent form	Yes	Pre-consent	No	The tool is built for use on a tablet in cars. The gamification targets only specific groups of individuals.
Our work	UI for cookie data visualisation.	Cookies	Yes	Post-consent	Yes	The tool focuses on visualising statistics of already accepted cookies. Revoking the consent for the cookies has not been implemented.

6. Conclusion and future work

In this paper, we presented an ontology^{4,5} for cookies and a KG-based tool⁶ for cookie information visualisation. The main goal of our solution focuses on easing users' comprehension of cookies and on raising awareness of cookie-based data sharing. The conducted user evaluation has shown that our approach to semantically representing and visualising cookies helps individuals understand the real nature of web cookies. In addition to empowering users with regard to their personal data sharing, we believe that this work helps to breach the gap between the Semantic Web and privacy domains with the help of the proposed cookie ontology.

Although the challenges of preserving individuals' privacy online and ensuring legally compliant cookie-based data sharing are far from being resolved, the rising interdisciplinary research between the legal, privacy and Semantic Web domains has already shown promising results. Ontologies such as ours can help to establish a reference model that eases domain experts' collaboration and semantically enriches the privacy-enhancing technologies and machine learning-based GDPR violation detection tools such as [5] that are being developed. Technologies such as SOLID [48] have been built to give individuals control over their personal data sharing online. While focusing on decentralisation of data, SOLID can still benefit from semantically representing cookies as a medium to request and receive consent online and can utilise the visualisation approach presented in this paper to support individuals' comprehension of data sharing. The results of the user evaluation have shown that individuals have limited and unclear understanding of the personal data about them that is collected and shared through cookies. However, the evaluation has also shown that our knowledge graph-based visualisation approach improves users' knowledge about cookies, privacy and the data sharing online.

Currently, our cookie visualisation tool is dependent on the Cookie Editor¹ extension and the information captured by it. Our future goal is to remove this dependency by extending the functionalities of our cookie visualisation tool (i.e., implement a cookie capture functionality). Another possible future direction is to extend the use case of our application such that not only it serves as a tool to communicate information, but it also allows users to act on it by offering them the possibility to manage cookies. On the semantic side, we have presented a novel ontology for cookies that can be extended for different domains and use cases. We believe that its reuse and extension will inspire

further collaboration between semantic and privacy experts. The uses of the KG for detecting security breaches and data-sharing patterns (within the cookies) can be explored as well.

Acknowledgements

This research is supported by the CampaNeo project funded by FFG (grant 873839) as well as the smashHit EU project funded under Horizon 2020 (grant 871477). We would like to thank Harshvardhan J. Pandit for sharing helpful insights on cookies, consent and GDPR.

References

- [1] J. Angulo, S. Fischer-Hübner, T. Pulls and E. Wästlund, Usable transparency with the data track, in: *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, 2015, pp. 1803–1808. doi:10.1145/2702613.2732701.
- [2] D.A. Audich, R. Dara and B. Nonnecke, Improving readability of online privacy policies through DOOP: A domain ontology for online privacy, *Digital* 1(4) (2021), 198–215. doi:10.3390/digital1040015.
- [3] A. Bechmann, Non-informed consent cultures: Privacy policies and app contracts on Facebook, *Journal of Media Business Studies* 11 (2015), 21–38. doi:10.1080/16522354.2014.11073574.
- [4] C. Bless, L. Dötlinger, M. Kaltschmid, M. Reiter, A. Kurteva, A.J. Roa-Valverde and A. Fensel, Raising awareness of data sharing consent through knowledge graph visualisation, in: *Studies on the Semantic Web*, 2021. doi:10.3233/ssw210034.
- [5] D. Bollinger, K. Kubicek, C. Cotrini and D. Basin, Automating cookie consent and {GDPR} violation detection, in: *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 2893–2910.
- [6] T. Brown et al., Design thinking, *Harvard business review* 86(6) (2008), 84.
- [7] J.M. Brunetti, S. Auer, R. García, J. Klímek and M. Nečaský, Formal linked data visualization model, in: *Proceedings of International Conference on Information Integration and Web-Based Applications & Services, IIWAS'13*, Association for Computing Machinery, New York, NY, USA, 2013, pp. 309–318. ISBN 9781450321136. doi:10.1145/2539150.2539162.
- [8] J. Cardoso and A. Sheth, in: *The Semantic Web and Its Applications*, 2006, pp. 3–33. ISBN 978-0-387-30239-3. doi:10.1007/978-0-387-34685-4_1.
- [9] T.R. Chhetri, A. Kurteva, J.G. Adigun and A. Fensel, Knowledge graph based hard drive failure prediction, *Sensors* 22(3) (2022), <https://www.mdpi.com/1424-8220/22/3/985>. doi:10.3390/s22030985.
- [10] T.R. Chhetri, A. Kurteva, R.J. DeLong, R. Hilscher, K. Korte and A. Fensel, Data protection by design tool for automated GDPR compliance verification based on semantically modeled informed consent, *Sensors* 22(7) (2022), <https://www.mdpi.com/1424-8220/22/7/2763>. doi:10.3390/s22072763.
- [11] S. Cox, H. Alani, H. Glaser and S. Harris, The semantic web as a semantic soup, in: *1st Workshop on Friend of a Friend, Social Networking and the Semantic Web*, 2004.
- [12] B. Custers, S. van der Hof and B. Schermer, Privacy expectations of social media users: The role of informed consent in privacy policies, *Policy & Internet* 6(3) (2014), 268–295. doi:10.1002/1944-2866.POI366.
- [13] S. de Lusignan, S. Shinneman, I. Yonova, J. van Vlymen, A. Elliot, F. Bolton, G. Smith and S. O'Brien, An ontology to improve transparency in case definition and increase case finding of infectious intestinal disease: Database study in English general practice, *JMIR Medical Informatics* 5 (2017), e34. doi:10.2196/medinform.7641.
- [14] A. Dimou, L. De Vocht, G. Van Grootel, L. Van Campe, J. Latour, E. Mannens and R. Van de Walle, Visualizing the information of a linked open data enabled research information system, *Procedia Computer Science* 33 (2014), 245–252, <https://www.sciencedirect.com/science/article/pii/S1877050914008291>. doi:10.1016/j.procs.2014.06.039.
- [15] O. Drozd and S. Kirrane, I agree: Customize your personal data processing with the CoRe user interface, in: *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2019, pp. 17–32. ISBN 9783030278120. doi:10.1007/978-3-030-27813-7_2.
- [16] O. Drozd and S. Kirrane, Privacy CURE: Consent comprehension made easy, 2020. ISBN 978-3-030-58200-5. doi:10.1007/978-3-030-58201-2_9.
- [17] European Parliament, Council of the European Union, Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), *Official Journal of the European Union L*, 119 (May 2016), <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.
- [18] European Parliament, Council of the European Union, Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications), 2002, <https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32002L0058>.
- [19] K. Fatema, E. Hadziselimovic, H.J. Pandit, C. Debruyne, D. Lewis and D. O'Sullivan, Compliance through informed consent: Semantic based consent permission and data management model, in: *PrivOn@ISWC*, 2017, http://ceur-ws.org/Vol-1951/PrivOn2017_paper_5.pdf.

- [20] D. Fensel, *Ontologies: A Silver Bullet for Knowledge Management and Electronic Commerce*, 2nd edn, Springer-Verlag, Berlin, Heidelberg, 2003. ISBN 3540003029. doi:10.1007/978-3-662-09083-1.
- [21] N. Freire and S. de Valk, Automated interpretability of linked data ontologies: An evaluation within the cultural heritage domain, in: *2019 IEEE International Conference on Big Data (Big Data)*, 2019, pp. 3072–3079. doi:10.1109/BigData47090.2019.9005491.
- [22] C.M. Gray, Y. Kou, B. Battles, J. Hoggatt and A.L. Toombs, The dark (patterns) side of UX design, in: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, Association for Computing Machinery, New York, NY, USA, 2018, pp. 1–14. ISBN 9781450356206. doi:10.1145/3173574.3174108.
- [23] P. Grünewald and T. Reisch, The trust gap: Social perceptions of privacy data for energy services in the United Kingdom, *Energy Research & Social Science* **68** (2020), 101534, <https://www.sciencedirect.com/science/article/pii/S2214629620301110>. doi:10.1016/j.erss.2020.101534.
- [24] V. Ha, K. Inkpen, F.A. Shaar and L. Hdeib, An examination of user perception and misconception of Internet cookies, in: *CHI'06 Extended Abstracts on Human Factors in Computing Systems, CHI EA'06*, Association for Computing Machinery, New York, NY, USA, 2006, pp. 833–838. ISBN 1595932984. doi:10.1145/1125451.1125615.
- [25] B. Heitmann and C. Hayes, An architecture and methodologies for federated, privacy-enabled personalisation on the web of data, *Semantic Web* (2011).
- [26] S. Human, H.J. Pandit, V.P. Morel, C. Santos, M. Degeling, A. Rossi, W. Botes, V. Jesus and I. Kamara, in: *Data Protection and Consenting Communication Mechanisms: Current Open Proposals and Challenges, International Workshop on Privacy Engineering – IWPE'22, Co-Located with 7th IEEE European Symposium on Security and Privacy*, Genoa, Italy, 6 June 2022, 2022.
- [27] S. Jegatheesan, Cookies invading our privacy for marketing advertising and security issues, 2013, arXiv preprint [arXiv:1305.2306](https://arxiv.org/abs/1305.2306).
- [28] Y. Jia, Y. Qi, H. Shang, R. Jiang and A. Li, A practical approach to constructing a knowledge graph for cybersecurity, *Engineering* **4**(1) (2018), 53–60. doi:10.1016/j.eng.2018.01.004.
- [29] R. Joergensen and I. Review, The unbearable lightness of user consent, *Internet Policy Review* **3** (2014). doi:10.14763/2014.4.330.
- [30] M.L. Jones, Cookies: A legacy of controversy, *Internet Histories* **4**(1) (2020), 87–104. doi:10.1080/24701475.2020.1725852.
- [31] S. Kirrane, S. Villata and M. d'Aquin, Privacy, security and policies: A review of problems and solutions with semantic web technologies, *Semantic Web* **9**(2) (2018), 153–161. doi:10.3233/SW-180289.
- [32] A. Kurteva, Making sense of consent with knowledge graphs, PhD thesis, Semantic Technology Institute (STI) Innsbruck, Department of Computer Science, University of Innsbruck. Available at <https://digital.obvsg.at/urn:urn:nbn:at:at-ubi:1-113241>.
- [33] A. Kurteva, T.R. Chhetri, H.J. Pandit and A. Fensel, Consent through the lens of semantics: State of the art survey and best practices, *Semantic Web* (2021), 1–27. doi:10.3233/SW-210438.
- [34] A. Mathur, M. Kshirsagar and J. Mayer, What makes a dark pattern... dark? Design attributes, normative considerations, and measurement methods, in: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–18. doi:10.1145/3411764.3445610.
- [35] C. Matte, N. Bielova and C. Santos, Do cookie banners respect my choice?: Measuring legal compliance of banners from IAB Europe's transparency and consent framework, in: *2020 IEEE Symposium on Security and Privacy (SP)*, 2020, pp. 791–809. doi:10.1109/SP40000.2020.00076.
- [36] A.D. Miyazaki, Online privacy and the disclosure of cookie use: Effects on consumer trust and anticipated patronage, *Journal of Public Policy & Marketing* **27**(1) (2008), 19–33. doi:10.1509/jppm.27.1.19.
- [37] N. Noy and D. McGuinness, Ontology development 101: A guide to creating your first ontology, *Knowledge Systems Laboratory* **32** (2001).
- [38] A. Oltramari, D. Piraviperumal, F. Schaub, S. Wilson, S. Cherivirala, T. Norton, N.Á. Russell, P. Story, J. Reidenberg and N. Sadeh, PrivOnto: A semantic framework for the analysis of privacy policies, *Semantic Web* **9** (2017), 1–19. doi:10.3233/SW-170283.
- [39] H.J. Pandit, C. Debruyne, D. O'Sullivan and D. Lewis, GConsent—a consent ontology based on the GDPR, in: *European Semantic Web Conference*, Springer, 2019, pp. 270–282. doi:10.1007/978-3-030-21348-0_18.
- [40] H.J. Pandit, A. Polleres, B. Bos, R. Brennan, B. Bruegger, F.J. Ekaputra, J.D. Fernández, R.G. Hamed, E. Kiesling, M. Lizar, E. Schlehahn, S. Steyskal and R. Wenning, Creating a vocabulary for data privacy, in: *On the Move to Meaningful Internet Systems: OTM 2019 Conferences*, H. Panetto, C. Debruyne, M. Hepp, D. Lewis, C.A. Ardagna and R. Meersman, eds, Springer International Publishing, Cham, 2019, pp. 714–730. ISBN 978-3-030-33246-4. doi:10.1007/978-3-030-33246-4_44.
- [41] E. Papadogiannakis, P. Papadopoulos, N. Kourtellis and E.P. Markatos, User tracking in the post-cookie era: How websites bypass GDPR consent to track users, in: *Proceedings of the Web Conference 2021, WWW'21*, Association for Computing Machinery, New York, NY, USA, 2021, pp. 2130–2141. ISBN 9781450383127. doi:10.1145/3442381.3450056.
- [42] A. Piplai, S. Mittal, A. Joshi, T. Finin, J. Holt and R. Zak, Creating cybersecurity knowledge graphs from malware after action reports, *IEEE Access* **8** (2020), 211691–211703. doi:10.1109/ACCESS.2020.3039234.
- [43] M. Poveda-Villalón, A. Gómez-Pérez and M.C. Suárez-Figueroa, Oops! (ontology pitfall scanner!): An on-line tool for ontology evaluation, *International Journal on Semantic Web and Information Systems (IJSWIS)* **10**(2) (2014), 7–34. doi:10.4018/ijswis.2014040102.
- [44] Y. Qi, R. Jiang, Y. Jia and A. Li, Attack analysis framework for cyber-attack and defense test platform, *Electronics* **9**(9) (2020), 1413. doi:10.3390/electronics9091413.
- [45] P. Raschke, A. Küpper, O. Drozd and S. Kirrane, Designing a GDPR-compliant and usable privacy dashboard, in: *Privacy and Identity Management. The Smart Revolution. Privacy and Identity 2017*. IFIP Advances in Information and Communication Technology, 2017. doi:10.1007/978-3-319-92925-5_14.
- [46] S.C. Rasmusen, M. Penz, S. Widauer, P. Nako, A. Kurteva, A. Roa-Valverde and A. Fensel, Raising consent awareness with gamification and knowledge graphs: An automotive use case, *International Journal on Semantic Web and Information Systems (IJSWIS)* **18**(1) (2022), 1–21.

- [47] A. Rossi and M. Palmirani, A visualization approach for adaptive consent in the European data protection framework, in: *2017 Conference for e-Democracy and Open Government (CeDEM)*, 2017, pp. 159–170. doi:[10.1109/CeDEM.2017.23](https://doi.org/10.1109/CeDEM.2017.23).
- [48] A.V. Sambra, E. Mansour, S. Hawke, M. Zereba, N. Greco, A. Ghanem, D. Zagidulin, A. Abounaga and T. Berners-Lee, Solid: A platform for decentralized social applications based on linked data, *MIT CSAIL & Qatar Computing Research Institute, Tech. Rep.* (2016).
- [49] I. Sanchez-Rola, M. Dell’Amico, P. Kotzias, D. Balzarotti, L. Bilge, P.-A. Vervier and I. Santos, in: *Can I Opt Out Yet?: GDPR and the Global Illusion of Cookie Control*, 2019, pp. 340–351. doi:[10.1145/3321705.3329806](https://doi.org/10.1145/3321705.3329806).
- [50] C. Santos, N. Bielova and C. Matte, Are cookie banners indeed compliant with the law? Deciphering EU legal requirements on consent and technical means to verify compliance of cookie banners, 2019, arXiv preprint [arXiv:1912.07144](https://arxiv.org/abs/1912.07144).
- [51] C. Santos, A. Rossi, L. Sanchez Chamorro, K. Bongard-Blanchy and R. Abu-Salma, Cookie banners, what’s the purpose? Analyzing cookie banner text through a legal lens, in: *Proceedings of the 20th Workshop on Workshop on Privacy in the Electronic Society, WPES’21*, Association for Computing Machinery, New York, NY, USA, 2021, pp. 187–194. ISBN 9781450385275. doi:[10.1145/3463676.3485611](https://doi.org/10.1145/3463676.3485611).
- [52] T.H. Soe, O.E. Nordberg, F. Guribye and M. Slavkovik, Circumvention by design – dark patterns in cookie consent for online news outlets, in: *Proceedings of the 11th Nordic Conference on Human–Computer Interaction: Shaping Experiences, Shaping Society*, Association for Computing Machinery, New York, NY, USA, 2020. ISBN 9781450375795. doi:[10.1145/3419249.3420132](https://doi.org/10.1145/3419249.3420132).
- [53] B. Steichen and B. Fu, Towards adaptive information visualization – a study of information visualization aids and the role of user cognitive style, *Frontiers in Artificial Intelligence* **2** (2019). doi:[10.3389/frai.2019.00022](https://doi.org/10.3389/frai.2019.00022).
- [54] A. Tauqeer, A. Kurteva, T.R. Chhetri, A. Ahmeti and A. Fensel, Automated GDPR contract compliance verification using knowledge graphs, *Information* **13**(10) (2022), 447. doi:[10.3390/info13100447](https://doi.org/10.3390/info13100447).
- [55] R. Tirtea, C. Castelluccia and D. Ikonou, Bittersweet cookies. Some security and privacy considerations, *European Union Agency for Network and Information Security-ENISA* (2011).
- [56] M. Trusov, L. Ma and Z. Jamal, Crumbs of the cookie: User profiling in customer-base analysis and behavioral targeting, *Marketing Science* **35**(3) (2016), 405–426. doi:[10.1287/mksc.2015.0956](https://doi.org/10.1287/mksc.2015.0956).
- [57] C. Utz, M. Degeling, S. Fahl, F. Schaub and T. Holz, (Un)informed consent: Studying GDPR consent notices in the field, in: *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS’19*, Association for Computing Machinery, New York, NY, USA, 2019, pp. 973–990. ISBN 9781450367479. doi:[10.1145/3319535.3354212](https://doi.org/10.1145/3319535.3354212).
- [58] C. Ware, *Information Visualization: Perception for Design*, Morgan Kaufmann, 2019. ISBN 0123814642. doi:[10.1016/C2009-0-62432-6](https://doi.org/10.1016/C2009-0-62432-6).
- [59] F. Zhang, N.J. Yuan, D. Lian, X. Xie and W.-Y. Ma, Collaborative knowledge base embedding for recommender systems, in: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD’16*, Association for Computing Machinery, New York, NY, USA, 2016, pp. 353–362. ISBN 9781450342322. doi:[10.1145/2939672.2939673](https://doi.org/10.1145/2939672.2939673).
- [60] K. Zhang and J. Liu, Review on the application of knowledge graph in cyber security assessment, *IOP Conference Series: Materials Science and Engineering* **768** (2020), 052103. doi:[10.1088/1757-899X/768/5/052103](https://doi.org/10.1088/1757-899X/768/5/052103).