

PrivComp-KG: Leveraging Knowledge Graph and Large Language Models for Privacy Policy Compliance Verification

Leon Garza¹, Lavanya Elluri², Anantaa Kotal³, Aritran Piplai¹, Deepti Gupta², and Anupam Joshi²

¹ University of Texas at El Paso, {lgarza3,miners.utep.edu, apiplai@utep.edu}

² Texas A&M University Central Texas, {elluri,d.gupta}@tamuct.edu

³ University of Maryland Baltimore County, {anantak1,joshi}@umbc.edu

Abstract. Data protection and privacy is becoming increasingly crucial in the digital era. Numerous companies depend on third-party vendors and service providers to carry out critical functions within their operations, encompassing tasks such as data handling and storage. However, this reliance introduces potential vulnerabilities, as these vendors’ security measures and practices may not always align with the standards expected by regulatory bodies. In response, federal and international organizations have enacted data protection laws, regulations, and guidelines that govern vendor management and data dissemination. Businesses are required, often under the penalty of law, to ensure compliance with the evolving regulatory rules. However, interpreting and implementing these regulations pose challenges due to their complexity. Regulatory documents are extensive, demanding significant effort for interpretation, while vendor-drafted privacy policies often lack the detail required for full legal compliance, leading to ambiguity. To ensure a concise interpretation of the regulatory requirements and compliance of organizational privacy policy with said regulations, we propose a Large Language Model (LLM) and Semantic Web based approach for privacy compliance. In this paper, we develop the novel Privacy Policy Compliance Verification Knowledge Graph, PrivComp-KG. The PrivComp-KG is designed to efficiently store and retrieve comprehensive information concerning privacy policies, regulatory frameworks, and domain-specific knowledge pertaining to the legal landscape of privacy. Using LLMs and Retrieval Augmented Generation, we identify the relevant sections in a privacy policy with corresponding regulatory rules. Our LLM-based retrieval system achieved a correctness score of 0.9. This information about individual privacy policies is populated into the PrivComp-KG. Combining this with the domain context and rules, the PrivComp-KG can be queried to check for compliance with privacy policies by each vendor against relevant policy regulations. We demonstrate the relevance of the PrivComp-KG, by verifying compliance of privacy policy documents for various organizations. This approach allows the policy writers to meaningfully understand the law’s requirements, identify gaps in their existing policies, and update based on the evolving regulations.

Keywords: Privacy Policy · Policy Compliance · Large Language Model
· Knowledge Graph.

1 Introduction

Every day, vast quantities of data are collected from diverse sources across the globe. This spans from individual interactions on social media to industrial sensor readings. These collections of data are essential as it provides valuable insights into consumer behaviors, market trends, and societal patterns. Businesses leverage this data to tailor their products and services, optimize operations, and gain a competitive edge in the market. Researchers analyze vast datasets to make scientific breakthroughs, improve healthcare outcomes, and address societal challenges. This evolution coincides with a pervasive integration of technology into our daily lives, resulting in the widespread collection of individual and aggregated data.

However, the surge in data collection has sparked concerns regarding data privacy. There is potential for misuse and unauthorized dissemination of consumers' private information by organizations collecting their data. Consequently, numerous data protection regulations, such as the European Union's General Data Protection Regulation (GDPR) [29], Payment Card Industry Data Security Standard (PCI DSS) [6], and the California Consumer Privacy Protection Act (CCPA) [1,2], have been established in response to public apprehension. Privacy Policy regulations impose strict rules on collecting, using, and managing personal data. These rules require companies to follow principles like minimizing data collection, limiting its use to specified purposes, obtaining user consent, and ensuring the data accuracy, security, and the company's accountability. The GDPR, for example, sets a thorough set of privacy and data protection rules for companies accessing the European Union (EU) users' data. Known for setting some of the strictest data privacy and security standards worldwide, the GDPR is a model for properly using personal data, focusing on safeguarding and legally processing individuals' information. As a result, businesses must reassess how they handle data, ensuring their approaches comply with GDPR guidelines to protect people's privacy while maximizing benefits with data in this digital era. Furthermore, not adhering to these regulations not only makes the organization more susceptible to data breaches, but also holds them liable to pay huge penalties.

In May 2023, the Irish Data Protection Commission (DPC) made a major move in the history of the GDPR by fining the American tech company Meta a record €1.2 billion [3]. This fine was the highest ever because Meta moved the personal data of European users to the United States without ensuring enough protection for this data. This step by the DPC is a crucial moment in data protection law, highlighting how seriously rules about international data transfer are enforced under the GDPR. In July 2021, Amazon Europe Core SARL was fined €746 million by the Luxembourg National Commission for Data Protection (CNDP), marking the most significant penalty for violating the GDPR. This

action followed a complaint from 10,000 individuals, organized by the French privacy group La Quadrature du Net, concerning Amazon’s data processing practices. The CNDP’s investigation revealed that Amazon’s advertising targeting system operated without obtaining proper consent, contravening GDPR’s stringent consent requirements, which demand clear communication and detailed explanations of personal data use, purpose, and usage. [25] In September 2023, TikTok faced a significant penalty from the Irish DPC, receiving a fine of EUR 345 million. This event marked one of the most considerable GDPR fines imposed on a social media platform, particularly highlighting issues around protecting children’s data privacy. The decision underscored the critical need for tech companies to prioritize the safety of young users online. [28] These instances underscore the critical need for businesses to understand privacy policy regulations and ensure that their data privacy policies are consistent with the regulations.

Businesses must thoroughly understand the nature, scope, and purpose of the data they collect, process, and store. Furthermore, this information must be precisely recorded in a privacy policy document that’s easy for users to access and understand. Writing a comprehensive privacy policy document is critical to building trust between data collectors and consumers. Developing a privacy policy that meets the extensive requirements of policy regulations is a significant challenge for companies, primarily because of the complexity of the regulation’s rules. Privacy policies are short and concise for user ease while complying with all relevant sections of the regulatory document. Furthermore, the legal landscape of data protection regulations is dynamic and evolving. The GDPR sets a high privacy and data protection standard, yet it’s subject to interpretation and ongoing adjustments by regulatory authorities. This ever-changing landscape necessitates that companies stay flexible and constantly monitor legal updates to ensure their privacy policies and practices align with compliance requirements. Additionally, the regulation applies to businesses operating within the EU and those outside the region if they process data from EU residents. This global reach complicates compliance efforts. Companies must navigate GDPR and any other local privacy laws that might apply. This complicates efforts to develop compliant privacy policies that satisfy all regulatory requirements. The motivation of this work is to help privacy policy writers efficiently develop a privacy policy in compliance with regulatory documents.

We propose an innovative framework that enables policy writers to identify relevant regulatory rules, detect shortcomings in existing policies, and effectively address compliance requirements. In this work, we develop a Privacy Policy Compliance Verification Knowledge Graph, PrivComp-KG, that is designed to collect and maintain information regarding policy and regulatory documents, and encapsulate domain knowledge. The inference rule engine is used to reason over the privacy policies and regulatory documents to verify compliance. The query engine is utilised to effectively gain this insight from the PrivComp-KG and can be utilised by policy writers to identify any gaps in their existing policies. To efficiently populate the PrivComp-KG with privacy policy documents, as well as

their relevance to regulatory sections, we use Retrieval Augmented Generation (RAG) to assess privacy policies' alignment with GDPR articles. This avoids the need for constant model fine-tuning with evolving privacy laws. RAG helps generate responses by utilizing chunks of GDPR articles, allowing us to identify specific segments that match privacy policies, ensuring a dynamic and comprehensive approach without requiring continual model updates for each policy change. We demonstrate the utility of the PrivComp-KG, by verifying compliance of privacy policy documents for various organizations in the OPP-115 dataset [33].

The structure of the paper is organized as follows. In Section 2 we describe relevant work in contemporary GDPR compliance efforts and LLM RAG methods for text retrieval. In Section 3, we describe our framework leveraging LLMs for RAG to align GDPR sections with privacy policies. This section also discusses the development of PrivComp-KG and illustrates how data retrieved from vendor policies are incorporated into the knowledge graph. In Section 4 we detail the utilisation of our framework to verify compliance of privacy policies in the OPP-115 dataset. Furthermore we describe the results from evaluation of the knowledge extraction and KG creation methods. The final section summarizes these discoveries and outlines prospective research avenues.

2 Background and Related Work

2.1 Privacy Policy Compliance

As data collection increases, public concern over privacy risks has intensified, leading to the implementation of stringent privacy regulations like the GDPR. These regulations are crucial for safeguarding privacy, yet ensuring automatic compliance remains a complex challenge. The GDPR marks a significant change in privacy laws, setting up strict rules to protect people's rights and how companies must handle personal data. The European Union introduced it, requiring strict adherence from any organization within the EU or dealing with data from EU citizens. Ignoring these rules can lead to heavy fines of up to 20 million Euros or 4% of a company's worldwide yearly income, depending on the seriousness of the violation. The GDPR applies to all organizations handling personal data, regardless of location, and replaces the older EU Data Protection Directive introduced in 1995, creating a unified privacy law for all EU countries. This means companies can now follow one set of privacy rules across the EU. [29] Furthermore, the GDPR's influence has reached beyond Europe, inspiring similar laws in other countries, including the United States. Given this move towards tighter privacy regulations globally, understanding and complying with the GDPR is crucial for software companies to meet current rules and prepare for any new privacy laws in the future.

In recent research [26], the concept of "Data Capsule" is introduced, automating compliance checks against privacy regulations in data processing. Individual data is associated with specific policies, ensuring adherence through residual policies and a new algorithm for effective policy derivation. This system

advances individual privacy protection, albeit focusing solely on data subject rights. Another study [17] proposes an approach to assess privacy policy alignment with GDPR Article 13 standards. By manually selecting 304 policies and developing a labeling system, the authors identify compliance issues, creating a web tool named AutoCompliance to simplify policy comprehension. However, this study overlooks broader GDPR coverage. In privacy policy research [16], the impact of GDPR on over 6,000 policies is analyzed, indicating significant revisions post-GDPR, particularly in EU policies. User experience improvements are noted, but confidence in vendor compliance remains uncertain. Leveraging ontology and text extraction techniques [14], vendors automate privacy policy compliance efforts, streamlining data protection measures. These advancements signify progress in managing privacy constraints, but comprehensive compliance assurance remains a challenge.

In our earlier research studies, we developed a foundational compliance knowledge graph to include various regulations and incorporated a selection of vendor privacy policy descriptions into the ontology without directly linking them to specific GDPR chapters or sections. Also, we correlated these documents with Cloud Security Alliance (CSA) controls to bridge gaps. In our past research [10,11], we identified relevant sections by extracting keywords and entities from the glossaries or appendices of regulations. We then identified the semantically similar keywords associated with GDPR regulation from the vendor privacy policies. Further, we checked for the semantic similarity between the summaries of the entire GDPR and the privacy policy document using a generic BERT abstractive summarize [9]. In another research work, we have incorporated the National Institute of Standards and Technology (NIST) 8228 [4] risk mitigation areas into the knowledge graph. [5, 8].

Expanding upon this groundwork, our current research seeks to align the extracted GDPR articles from vendor privacy policies, pinpointing any previously overlooked articles. This effort is designed to assist vendors in refining their policy documents. Given the extensive nature of these regulations, our focus is primarily on GDPR compliance, recognizing its significance for vendors handling data from EU users. The traditional approach often necessitates manual review to ensure compliance. However, our methodology proposes a more efficient solution for identifying and addressing gaps in vendor privacy policies, reducing the need for human intervention.

2.2 Information Extraction using LLMs

Although mostly statistical and rule-based methods have been used for information extraction [22, 23], recent advancements in LLMs have opened up new methodologies. LLMs have been extensively used for information extraction across various domains. For information extraction tasks, LLMs have been utilized for generating structured entities and relationships circumventing the need to use supervised models [15, 20, 34]. However, to find more success in specific domains, LLMs have been fine-tuned to perform the task of information extraction. For example, in the case of scientific data extraction [7] and agriculture data




Fig. 1: Overall Framework for Building and Querying GDPR Vendor Policy Management Knowledge Graph

extraction [21], fine-tuned LLMs have been used. LLMs have also been used for improving annotations in the medical domain, by periodically fine-tuning based on human feedback [13]. In the domain of cybersecurity, LLMs have also been used for information combination and extraction [18, 19]. However, training and fine-tuning an LLM is computationally expensive and there is little guarantee that the model will not suffer from hallucinations. In our approach, we have utilized the power of RAG to limit the possibilities of hallucinations and avoid the cost of fine-tuning for our specific application scenario.

3 Methodology

The Privacy Policy Compliance Verification Knowledge Graph (PrivComp-KG) formalizes GDPR rules and guidelines using Semantic Web technologies. It facilitates automated compliance checking, enhances transparency, and supports granular consent management. The Knowledge Graph allows for swift cross-

referencing of regulatory requirements and vendor privacy policies, enabling efficient management of vendor data and adherence to regulations. Compliance inference using SWRL (Semantic Web Rule Language) rules [27] enriches the understanding of privacy policies, facilitating dynamic compliance and reasoning over gaps in policies. The PrivComp-KG is populated with relevant privacy policy properties by leveraging LLMs to assess privacy policies’ alignment with regulatory articles, such as GDPR. We employ Retrieval-Augmented Generation (RAG) to mitigate LLMs’ tendency for hallucinations when confronted with unfamiliar queries, dynamically generating responses without continual fine-tuning. The overall framework for PrivComp-KG creation and population using LLMs and end user querying is demonstrated in Figure 1.

3.1 Knowledge Extraction using LLM

LLMs are increasingly valued for their deep understanding of natural language. The vector representations of each piece of text written in natural language, have a deeper contextual meaning in the scope of LLMs.

To understand privacy policies, we harness the capabilities of LLMs to assess the alignment of a privacy policy with GDPR articles. Previous efforts have focused on fine-tuning LLMs to identify similar GDPR entities corresponding to a specific privacy policy [9]. However, both LLMs in general and fine-tuned LLMs specifically are prone to hallucinations when confronted with queries relating to unfamiliar domains. An increasingly popular method to address this issue is RAG. We opt for RAG due to the dynamic nature of the domain. Given that data privacy laws can be evolving, characterized by the emergence of new regulations, RAG aids in response generation without the need for continual model fine-tuning for every update in privacy laws and regulations.

The core components of an LLM consist of (i) a query or prompt, denoted as P , and (ii) a response, represented by R . In the context of RAG-enabled LLMs, the generation of R relies on a set of documents, denoted as D . We utilize RAG not only to produce responses for prompts inquiring about the relationship between a privacy policy and a GDPR article but also to identify the specific segments of a GDPR article that align with a privacy policy.

We segment each GDPR article into chunks and integrate them into a vector store. Each segment of the GDPR article is assigned a representation within the vector store. Our vector store contains 430 chunks of GDPR articles, with each article further divided into multiple segments. Typically, each paragraph of a GDPR article corresponds to one vector within our vector store. An example of our prompt P is illustrated in Figure 2.

As part of the metadata for the chunks inserted into the vector store, we include the specific GDPR article from which each chunk was extracted. A typical P for this system might be "Which GDPR article does this privacy policy relate to?" followed by the privacy policy itself. The retriever then (i) generates R based on the GDPR articles and (ii) provides a list of articles used to synthesize R , along with their corresponding similarity scores.




Fig. 2: Example of our model’s response generation on a privacy policy sample

In Figure 2, we observe an example of the model’s performance using an excerpt from Microsoft’s XBOX privacy policy. Our vector store is constructed by inserting chunks of GDPR articles, which are then utilized by our LLM (LLama-7B) to generate R . In the example, we highlight Article 21, which achieved the highest similarity score with P . During our experiments, we establish a threshold for the similarity score and list all articles used in generating R .

This process creates a comprehensive list of articles that correlate with a given privacy policy. Our system supports dynamic updates to privacy policies, as model retraining is unnecessary when policies are amended or added. The knowledge derived from privacy policies and their corresponding articles can be incorporated into our PrivComp-KG. In the following sections, we elaborate on how the reasoning capabilities of a knowledge graph can be leveraged to derive valuable insights from privacy policies and the associated GDPR articles.

3.2 PrivComp-KG: Privacy Policy Compliance Verification Knowledge Graph

Leveraging Semantic Web technologies can streamline regulatory compliance. Using the Semantic Web for privacy policy compliance offers advantages such as standardization and machine readability. It enables automated compliance checking by representing policies in a format that software tools can understand and analyze. This approach enhances transparency by allowing policy writers to easily comprehend data privacy requirements and gaps in their existing policies. Furthermore, the semantic representations facilitate granular consent management, tailoring policies to specific regulations.

We utilized the semantic web languages Resource Description Framework (RDF) [31] and Web Ontology Language (OWL) [30] to capture and formalize




Fig. 3: PrivComp-KG: Privacy Policy Compliance Verification Knowledge Graph along with Instances of the classes

the rules and guidelines outlined in GDPR and Vendor policy documents. We developed the novel Privacy Compliance Verification KG, PrivComp-KG. It is designed to be in the public domain and can be adopted quickly and easily by vendors who are seeking to adhere to these regulations. The ontology is also platform-independent and can be integrated with the latest data protection regulations and many other data regulation entities.

RDF enhances the structuring of knowledge on the web, simplifying the retrieval of domain-specific information for vendors. PrivComp-KG is integrated with our existing Reference Document Knowledge Graph [9], allowing for the swift and effective cross-referencing of regulatory requirements and vendor privacy policies. As illustrated in Figure 3, this high-level knowledge graph manages GDPR rules and vendor policy extracted results. This knowledge graph is specifically designed to accommodate any applicable regulations for various vendors or companies based on the types of data they collect. Our focus on GDPR, a pivotal regulation with extensive stipulations, forms a critical component of this research methodology. Utilizing Protege software [24,32], a free, open-source platform for ontology editing and reasoning, we constructed and managed our ontology. This structured and standardized approach not only facilitates the management of vendor data but also ensures adherence to regulations, thereby safeguarding customer data and privacy. Our knowledge graph is hosted in a public space, providing accessible user interaction.

In our earlier research, described in [9], we could only retrieve high-level entities and didn't examine the specific rules or articles related to privacy poli-

cies. Additionally, our previous version of the knowledge graph categorized most GDPR rules as classes, which made it challenging to compare results across different sections of the GDPR. Structuring these as instances improved our ability to identify what was missing quickly or had been extracted. We used the insights from the Section 3.1 to update the data properties of the *Provider* class. With this research, we’ve successfully refined our knowledge graph to manage and search GDPR rules flexibly and compare them with those extracted from vendor policy documents. We designed this knowledge graph to include a Regulations class that branches into various regulations like PCI-DSS, HIPAA, CCPA, etc. This study extracted relevant chapters, articles, and obligations from the GDPR and stored them as instances in *GDPR_Chapter*, *GDPR_Articles*, and *GDPR_Obligations*.

3.3 Regulatory Obligations:

Beyond the curation of Provider Privacy Policies and Policy Regulations, the key insight in PrivComp-KG is drawn from the Regulatory Obligations instances. Every regulation describes based on the role of the data actor, the specific set of rules that apply to the data actor. For example, a provider, i.e. an organization offering a specific service, may collect data to support that service. However, GDPR clearly states in its provider guidelines, the specific measures that the providers need to take to protect data privacy. Additionally, some general rules with regards to the ethics and responsibility of data collection apply to the providers. This knowledge about the specific role and application of GDPR sections is encoded in the *GDPR_Obligations* class. Specifically, there are 5 instances of the *GDPR_Obligations*, each describing a specific role:

- *Consumer_Obligations* Consumers must inform the supervisory authority and the data subject about any personal data breaches. Additionally, the consumer must conduct a Data Protection Impact Assessment (DPIA), consult with the supervisory authority before processing if the DPIA indicates a high risk, and appoint a Data Protection Officer when processing personal data on a large scale.
- *Common_Obligations* Rules that apply to both consumers and providers, who are responsible for ensuring compliance.
- *Data_Subject_Obligations* Consumers must inform data subjects about the duration for which, or why, data will be retained upon collection. Consequently, if data subjects request the removal of their data, and it is no longer necessary for the purposes for which it was collected, it must be erased.
- *General_Obligations* GDPR mentions generic rules that are not specific to any role but apply to all the data processing activities in general.
- *Provider_Obligations* Provider is primarily responsible for assisting consumers in the event of data breaches and processing data in accordance with consumer directives. Additionally, the Provider must maintain comprehensive records of all data processing activities and ensure robust data security measures are in place to protect consumer information.

3.4 Compliance and Regulatory Properties:

The PrivComp-KG supports object properties that links classes to support compliance reasoning and verification.

- *hasRegulation*: Regulations identified in vendor privacy policies are stored as instances within the *Regulation* class and linked to the *Providers* class.
- *compliesWithSection*: Using the knowledge extraction tool, as described in Section 3.1, the regulatory articles that individual provider privacy policies complies with are identified. This knowledge is populated into the PrivComp-KG using the "compliesWithSection" relation between *Providers* and *GDPR_Articles*.
- *requiresComplianceWith*: Every vendor handling EU user data must adhere to relevant GDPR requirements. To enforce this, we link the *Providers* class to the *GDPR_Chapters* through this object property.
- *partOfChapter*: Since a chapter can encompass multiple articles, we connect the *GDPR_Articles* to the *GDPR_Chapters* class to reflect this relationship.

Additionally, to organize the findings from privacy policy analyses, we established several data properties:

- *hasChapterIndex*: This property stores the indices of chapters identified in a policy document.
- *hasSectionIndex*, *hasSectionText*, and *hasSectionURL*: These properties are essential for recording the sections extracted from the documents, including descriptions and URLs for easy reference.

3.5 Compliance Inference

SWRL rules enhance reasoning capabilities in Semantic Web applications by allowing the specification of logical rules that define relationships and infer new information from existing data. The PrivComp-KG supports SWRL rules that infers necessary rules from regulatory articles based on the role of the data actor and the role based obligations, as described in Section 3.3.

For example, for the privacy policy of a provider that wishes to collect and utilise data, the following SWRL rules can infer the obligatory rules from GDPR.

```
S1: Cloud_Providers(?cloud_provider) ^ GDPR_Articles(?gdpr_article)
  ^ definesObligationsFor(?gdpr_article, Provider_Obligations)
  -> requiresComplianceWith(?cloud_provider, ?gdpr_article)

S2: Cloud_Providers(?cloud_provider) ^ GDPR_Articles(?gdpr_article)
  ^ definesObligationsFor(?gdpr_article, Common_Obligations)
  -> requiresComplianceWith(?cloud_provider, ?gdpr_article)
```

By applying these SWRL rules, PrivComp-KG can derive detect inconsistencies in existing policies, and make logical deductions about compliance. This reasoning process enriches the understanding and interpretation of privacy policy data, facilitating a dynamic approach towards privacy policy compliance, more advanced semantic querying and data integration. After processing with the reasoner, users can efficiently compare the required rules against those extracted, updating any lacking areas in the privacy policy to ensure it is current and compliant.

4 Experimental Results

4.1 Dataset

The OPP-115 dataset [33] is a comprehensive collection of privacy policies from various online platforms, consisting of over 115 privacy policies. It encompasses a wide range of websites and services, including social media platforms, e-commerce sites, and mobile applications. The dataset is structured and annotated, making it suitable for privacy policy analysis and evaluating language models. It includes the categorisation of data collection, usage, sharing, and retention practices outlined in the privacy policies. In this work, we use the OPP-115 dataset to demonstrate the utility of our model and evaluate the performance of the LLMs. Our LLM based method is used to identify relevant regulatory articles for each provider in the dataset and then populated into PrivComp-KG. The inferential engine in PrivComp-KG reasons over the gaps in these provider policies. The results are made available to end users using a query engine. The results from our evaluation methods and query engine are described in subsequent sections.

4.2 Evaluation of LLM-guided extraction

For our experiments, we use Llama-7B LLM, and *chromaDB* as our vector store. The knowledge extracted from the RAG-enabled LLM focuses on GDPR articles corresponding to specific portions of a privacy policy. We evaluate this knowledge in two ways: (i) assessing the quality of RAG-LLM generated responses and (ii) determining the accuracy of the retrieved article numbers by the LLM.

As RAG is integral to the LLM, we assess the effectiveness of R generated by our model given P . A common metric for evaluating RAG-generated responses is RAGAS [12]. Our model’s

Threshold Used	Correctness Score
0.9	0.66
1.0	0.74
1.1	0.82
1.2	0.84
1.3	0.89
1.4	0.88
1.5	0.9

Table 1: Correctness score for each threshold




Fig. 4: RAGAS [12] score evaluating the quality of privacy policy "Answer" generation through our LLM, in comparison with human-generated "Answers". A score of '1' means perfect alignment.

responses are compared with those generated by humans. Figure 4 illustrates the correctness scores for all questions. Questions with low scores (<0.5) required more specific information regarding the privacy policies. Since our LLM only accesses GDPR articles, it performs well in questions regarding the similarity of a privacy policy to a GDPR article but struggles with overly specific questions about the policy itself.

The second metric utilizes the OPP-15 dataset. Since the OPP-15 dataset maps privacy policies to GDPR articles through an intermediate category, we leverage it to assess the accuracy of the articles retrieved during R generation. We pre-process the dataset to eliminate portions of a privacy policy that are too small. RAG offers a similarity score function, which we apply a threshold to, selecting only the chunks that fall below a certain threshold. Table 1 presents the accuracy of correctness for various thresholds.

4.3 Knowledge Graph Inferences and Queries

The PrivComp-KG is populated with the provider privacy policies from the OPP-115 dataset. Furthermore, the identified regulatory sections of relevance to the privacy policies are populated with "compliesWithSection" object property. Finally, using the reasoner, PrivComp-KG automatically updates the GDPR articles that are mandatory for the given provider using the "requiresComplianceWith" object property.

Fig 5 demonstrates the results for the provider "Lids.com". Leveraging the LLM results, the PrivComp-KG is populated with the GDPR articles that "Lids.com" has compliance with, as demonstrated in Fig 5 (a). Using the KG rea-

Property assertions: lids.com	Property assertions: lids.com
Object property assertions +	
compliesWithSection Art.4	requiresComplianceWith Art.79
compliesWithSection Art.5	requiresComplianceWith Art.57
compliesWithSection Art.23	requiresComplianceWith Art.35
compliesWithSection Art.45	requiresComplianceWith Art.58
compliesWithSection Art.34	requiresComplianceWith Art.36
compliesWithSection Art.88	requiresComplianceWith Art.77
compliesWithSection Art.14	requiresComplianceWith Art.11
compliesWithSection Art.25	requiresComplianceWith Art.55
compliesWithSection Art.13	requiresComplianceWith Art.56
compliesWithSection Art.46	requiresComplianceWith Art.78
	requiresComplianceWith Art.39

(a) GDPR articles with which the privacy policy currently complies

(b) GDPR articles with which the privacy policy needs to comply

Fig. 5: Compliance results for Lids.com from PrivComp-KG

soner, PrivComp-KG also populates the GDPR articles that require compliance for "Lids.com", as demonstrated in Fig 5 (b). These information are critical to policy writers of providers, like "Lids.com", in identifying gaps in their existing policies. To query the missing articles for a provider that need to be addressed, we provide the following SPARQL query:

```
SELECT * WHERE cc:lids.com
cc:requiresComplianceWith ?x MINUS
cc:lids.com cc:compliesWithSection ?x
```

For "Lids.com", the PrivComp-KG return 40 articles that has been identified as mandatory according to GDPR, but missing from the existing privacy policy. This information helps easily identify the gaps in the privacy policy and effectively identify the guidelines that need to be addressed.

5 Conclusion and Future Work

In the digital age, safeguarding data protection and privacy has become paramount. To effectively manage their operations, many companies enlist the support of third-party vendors and service providers. These partners play critical roles in various aspects of business functions, such as data handling, storage, and processing. Entrusting sensitive data to third parties necessitates a rigorous approach to ensure comprehensive data protection and privacy. Companies must establish robust protocols and contractual agreements to safeguard the integrity and confidentiality of the information shared with these external entities. Most existing research works focus only on specific sections of the regulation, which won't be helpful for organizations. In this research, we leverage LLM and Semantic Web technologies to verify compliance of privacy policy documents with policy

regulations, like GDPR. This work also proposed the novel Privacy Policy Compliance Verification Knowledge Graph, PrivComp-KG, for storing and retrieving comprehensive information to maintain privacy policies. This proposed research enhances the readability of privacy policies and also promotes transparency, empowers consumers, strengthens regulatory compliance, and ultimately fosters trust in the digital ecosystem. We aim to further assist companies by enriching our PrivComp-KG with the requisite U.S. federal and state-level data protection regulations, facilitating rapid cross-referencing across this comprehensive legislative framework.

Disclosure of Interests. Authors have no competing interests.

Supplemental Material Statement:

- The PrivComp-KG is available from:
<https://github.com/<anonaauthor>/PrivComp-KG.git>
- Datasets are available from: <https://www.usableprivacy.org/data> [33]
- Source code for populating PrivComp-KG can be accessed from:
<https://github.com/<anonaauthor>/PrivComp-KG.git>

References

1. California consumer privacy act of 2018. Tech. rep., State of California (2018)
2. California consumer privacy act of 2020. Tech. rep., State of California (2020)
3. Board, E.D.P.: Meta. https://www.edpb.europa.eu/news/news/2023/12-billion-euro-fine-facebook-result-edpb-binding-decision_en
4. Boeckl, K., Boeckl, K., Fagan, M., Fisher, W., Lefkowitz, N., Megas, K.N., Nadeau, E., O'Rourke, D.G., Piccarreta, B., Scarfone, K.: Considerations for managing Internet of Things (IoT) cybersecurity and privacy risks. US Department of Commerce, National Institute of Standards and Technology ... (2019)
5. Boeckl, K., Boeckl, K., Fagan, M., Fisher, W., Lefkowitz, N., Megas, K.N., Nadeau, E., O'Rourke, D.G., Piccarreta, B., Scarfone, K.: Considerations for managing Internet of Things (IoT) cybersecurity and privacy risks. US Department of Commerce, National Institute of Standards and Technology ... (2019)
6. Council, P.S.: Pci mobile payment acceptance security guidelines. https://listings.pcisecuritystandards.org/documents/PCI_Mobile_Payment_Acceptance_Security_Guidelines_for_Developers_v2_0.pdf (04 2024)
7. Dagdelen, J., Dunn, A., Lee, S., Walker, N., Rosen, A.S., Ceder, G., Persson, K.A., Jain, A.: Structured information extraction from scientific text with large language models. *Nature Communications* **15**(1), 1418 (2024)
8. Echenim, K.U., Elluri, L., Joshi, K.P.: Ensuring privacy policy compliance of wearables with iot regulations. In: 2023 5th IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA). pp. 247–256. IEEE (2023)
9. Elluri, L., Chukkapalli, S.S.L., Joshi, K.P., Finin, T., Joshi, A.: A bert based approach to measure web services policies compliance with gdpr. *IEEE Access* **9**, 148004–148016 (2021)

10. Elluri, L., Joshi, K.P., Kotal, A.: Measuring semantic similarity across eu gdpr regulation and cloud privacy policies. In: 2020 IEEE International Conference on Big Data (Big Data). pp. 3963–3978. IEEE (2020)
11. Elluri, L., Nagar, A., Joshi, K.P.: An integrated knowledge graph to automate gdpr and pci dss compliance. In: 2018 IEEE International Conference on Big Data (Big Data). pp. 1266–1271. IEEE (2018)
12. Es, S., James, J., Espinosa-Anke, L., Schockaert, S.: Ragas: Automated evaluation of retrieval augmented generation. arXiv preprint arXiv:2309.15217 (2023)
13. Goel, A., Gueta, A., Gilon, O., Liu, C., Erell, S., Nguyen, L.H., Hao, X., Jaber, B., Reddy, S., Kartha, R., et al.: Llms accelerate annotation for medical information extraction. In: Machine Learning for Health (ML4H). pp. 82–100. PMLR (2023)
14. Hu, Y.J., Guo, H.Y., Lin, G.D.: Semantic enforcement of privacy protection policies via the combination of ontologies and rules. In: 2008 IEEE International Conference on Sensor Networks, Ubiquitous, and Trustworthy Computing (SUTC 2008). pp. 400–407. IEEE (2008)
15. Li, Z., Zeng, Y., Zuo, Y., Ren, W., Liu, W., Su, M., Guo, Y., Liu, Y., Li, X., Hu, Z., et al.: Knowcoder: Coding structured knowledge into llms for universal information extraction. arXiv preprint arXiv:2403.07969 (2024)
16. Linden, T., Khandelwal, R., Harkous, H., Fawaz, K.: The privacy policy landscape after the gdpr. arXiv preprint arXiv:1809.08396 (2018)
17. Liu, S., Zhao, B., Guo, R., Meng, G., Zhang, F., Zhang, M.: Have you been properly notified? automatic compliance analysis of privacy policy text with gdpr article 13. In: Proceedings of the Web Conference 2021. pp. 2154–2164 (2021)
18. Liu, Z., Shi, J., Buford, J.F.: Cyberbench: A multi-task benchmark for evaluating large language models in cybersecurity
19. Mitra, S., Neupane, S., Chakraborty, T., Mittal, S., Piplai, A., Gaur, M., Rahimi, S.: Localintel: Generating organizational threat intelligence from global and local cyber knowledge. arXiv preprint arXiv:2401.10036 (2024)
20. Peng, L., Wang, Z., Yao, F., Wang, Z., Shang, J.: Metaie: Distilling a meta model from llm for all kinds of information extraction tasks. arXiv preprint arXiv:2404.00457 (2024)
21. Peng, R., Liu, K., Yang, P., Yuan, Z., Li, S.: Embedding-based retrieval with llm for effective agriculture information extracting from unstructured data. arXiv preprint arXiv:2308.03107 (2023)
22. Pingle, A., Piplai, A., Mittal, S., Joshi, A., Holt, J., Zak, R.: Relext: Relation extraction using deep learning approaches for cybersecurity knowledge graph improvement. In: Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining. pp. 879–886 (2019)
23. Piplai, A., Mittal, S., Joshi, A., Finin, T., Holt, J., Zak, R.: Creating cybersecurity knowledge graphs from malware after action reports. IEEE Access **8**, 211691–211703 (2020)
24. Protege: Protege tool. <http://protege.stanford.edu> (2020). <https://doi.org/10.000/55555>, accessed: March 20, 2024
25. Review, N.L.: Luxembourg amazon gdpr violations. <https://natlawreview.com/article/luxembourg-dpa-fines-amazon-746-million-euros-gdpr-violations>
26. Springer: Data capsule: A new paradigm for automatic compliance with data privacy regulations (2019)
27. SWRL, A.: Semantic web rule language combining owl and ruleml. W3C Member Submission (May 21, 2004), <http://www.w3.org/Submission/SWRL/> (last visited March 2011) (2004)

28. Times, T.I.: Tiktok fined €345m by ireland’s data regulator for violating children’s privacy. <https://www.irishtimes.com/technology/big-tech/2023/09/15/tiktok-fined-345m-by-irelands-data-regulator-for-violating-childrens-privacy/>
29. Union, E.: General data protection regulation. <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng> (04 2024)
30. W3C: Web Ontology Language (10 February, 2004), <https://www.w3.org/TR/owl-features/>, accessed: August 23, 2023
31. W3C: Resource Description Framework (15 March, 2014), <https://www.w3.org/RDF/>, accessed: August 23, 2023
32. W3C: Protégé (software). [https://en.wikipedia.org/wiki/Prot%C3%A9g%C3%A9_\(software\)](https://en.wikipedia.org/wiki/Prot%C3%A9g%C3%A9_(software)) (2022, March 16), accessed: March 20, 2024
33. Wilson, S., Schaub, F., Dara, A.A., Liu, F., Cherivirala, S., Leon, P.G., Andersen, M.S., Zimmeck, S., Sathyendra, K.M., Russell, N.C., et al.: The creation and analysis of a website privacy policy corpus. In: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 1330–1340 (2016)
34. Xu, D., Chen, W., Peng, W., Zhang, C., Xu, T., Zhao, X., Wu, X., Zheng, Y., Chen, E.: Large language models for generative information extraction: A survey. arXiv preprint arXiv:2312.17617 (2023)