Unifying Large Language Models and Knowledge Graphs for efficient Regulatory Information Retrieval and Answer Generation

Kishore Vanapalli
Carleton University, Ottawa, Canada
kishorevanapalli@cmail.carleton.ca

Aravind Kilaru
Carleton University, Ottawa, Canada
aravindkilaru@cmail.carleton.ca

Omair Shafiq
Carleton University, Ottawa, Canada
omairshafiq@cunet.carleton.ca

Shahzad Khan
Gnowit Inc., Ottawa, Canada
shahzad@gnowit.com

Abstract

In a rapidly changing socio-economic landscape, regulatory documents play a pivotal role in shaping responses to emerging challenges. An efficient regulatory document monitoring system is crucial for addressing the complexities of a dynamically evolving world, enabling prompt crisis response, simplifying compliance, and empowering data-driven decision-making. In this work, we present a novel comprehensive analytical framework, PolicyInsight, which is based on a specialized regulatory data model and state-of-the-art NLP techniques of Large Language Models (LLMs) and Knowledge Graphs to derive timely insights, facilitating data-driven decision-making and fostering a more transparent and informed governance ecosystem for regulators, businesses, and citizens.

1 Introduction

1.1 Problem Statement

Regulatory policy monitoring (Waterman and Wood, 1993) refers to the systematic process of observing, tracking, and analyzing the policies and regulations established by regulatory bodies. The primary goal is to stay informed about any changes, updates, or new developments in regulatory policies that may affect various sectors, industries, or the general public. This monitoring process involves continuous observation, change detection, impact analysis, and compliance monitoring.

1.2 Importance

Monitoring and tracking regulatory policies are highly important for businesses for several reasons, such as regulatory compliance, risk mitigation, strategic planning, operational efficiency, and market intelligence.

1.3 Difficulty

However, regulatory policy monitoring can be a challenging task due to various factors which include frequent policy changes, diverse regulatory frameworks, legislative complexity, lack of centralized information, data security and privacy challenges, and technological and automation challenges.

1.4 Solution

In this work, our objective was to develop an efficient and comprehensive regulatory document monitoring framework with the following features:
Real-time monitoring: The framework involves real-time monitoring of regulatory policy documents, ensuring that the information is always up-to-date.
Adaptability to Changes: With a novel policy data model, the system is designed to seamlessly adapt to changes in the structure or content of policy documents. It can dynamically adjust to modifications in document formats, new policy sections, or alterations in the way information is presented.
Intelligent Analytical Insights: State-of-the-art NLP techniques and LLMs (Pouyanfar et al., 2018; Zhou et al., 2020) are leveraged for better understanding and categorization of policy content and derive change detection and impact analysis.
Responsive User Interface: The user interface of the monitoring system is responsive and user-friendly. It allows users to interact with the data that are interested in a dynamic manner, facilitating efficient exploration, analysis, and tracking of policy updates.
Automated Alerts and Notifications: The dynamic approach includes the implementation of automated alert systems. Users can receive notifications in real time when significant policy changes occur, allowing for prompt response and analysis.
Scalability and performance: The system is

Proceedings of the 31st International Conference on Computational Linguistics, pages 22–30
January 19–24, 2025.

designed to scale efficiently, accommodating an increasing volume of policy documents and users. Performance optimization is a key aspect to ensure that the dynamic monitoring process remains efficient even as the dataset grows.

1.5 Scope

The scope of the paper is limited to the development of the following foundational features: Design and development of a novel and efficient data model (Devedžić, 1999) to organize, store, access and efficiently manage policy data. Using this data model, relationships between different policies or different versions of the same policy can be easily derived and utilized. Also, the new data model with the aid of relationships and constraints helps to derive key insights from the underlying policy data. Development of an advanced, intelligent and configuration driven Policy Monitoring Component which can collect, extract and store various policy data. Development of a sophisticated Policy Analytical System based on LLMs and Knowledge Graphs to achieve policy deduplication, policy impact analysis and policy change predictions. The policy data model and Knowledge Graph populated data, Cypher queries, sample LLM prompts, and evaluation results are shared in the project GitHub page[1].

2 Literature Review

Pan et al. (2024) proposed a futuristic roadmap for the unification of LLMs and Knowledge Graphs (KGs) to simultaneously leverage their advantages and proposed a roadmap which consists of three broad frameworks, specifically, 1) KG-enhanced LLMs, which leverage KGs during the pre-training and inference phases of LLMs, or for improving understanding of the knowledge gained by LLMs; 2) LLM-augmented KGs, that incorporates LLMs for different KG tasks such as embedding, graph-to-text generation, construction, completion and question answering; and 3) Synergized LLMs + KGs, in which LLMs and KGs, both provide equal contributions and work in a mutually beneficial way to improve both LLMs and KGs for bidirectional reasoning driven by both data and knowledge. Overall, the authors highlighted how LLMs and KGs complement each other in effectively addressing common challenges in several downstream tasks like Question-Answering, Hallucination detection and Reasoning.

Knowledge Graphs (KGs), which represent semantic relationships between entities, have shown significant relevance for NLP. Schneider et al. (2022) presented the results of an extensive survey, offering a multi-perspective review of tasks, research types, and contributions. It provides a structured overview of the research landscape, including a broad categorization of tasks, a summary of findings, and highlighted directions for future work after systematically analyzing over five hundred papers on Knowledge Graphs in NLP. The findings indicate that a wide range of tasks related to KGs in NLP have been studied across various domains, including emerging topics like knowledge graph embedding and augmented language models.

In the survey paper on Knowledge Graphs (KGs), Ji et al. (2022) provided a comprehensive review of knowledge graph covering overall research topics about 1) knowledge graph representation learning, 2) knowledge acquisition and completion, 3) temporal knowledge graphs, and 4) knowledge-aware applications, and summarize recent breakthroughs and perspective directions to facilitate future research. However, the paper fails to address some key aspects of KGs particularly while building and maintaining KGs and the way to overcome such challenges.

The survey by Abu-Salih (2021) is pioneering in providing a comprehensive definition of a domain-specific Knowledge Graph. Additionally, the paper conducts an extensive review of state-of-the-art approaches from academic works across seven domains of knowledge. However, it remains unclear why the discussed challenges cannot be generalized to domain-agnostic KGs, making it difficult to apply the solutions universally to any Knowledge Graph.

Dessì et al. (2021) introduced an innovative architecture that leverages natural language processing and machine learning (ML) techniques to extract entities and relationships from research publications, integrating them into a large-scale knowledge graph. However, as the paper notes, there are some limitations to the developed pipeline. For example, the current version does not fully utilize the semantic characterization of research entities to verify the resulting triples.

Johann Höchtl and Schöllhammer (2016) seeks

[1] Project GitHub page: `https://github.com/Kishorevb/policyinsight`

to bridge the gap between e-governance and public administration theories, moving beyond the predominantly service delivery-focused approach in much of e-government research. By utilizing the policy cycle as a model for policy processes and development, the article presents an innovative perspective on policy decision-making through the use of ICT and Big Data. It explores the delicate balance between the socially beneficial uses of Big Data and the potential harm to privacy and other values. This raises complex questions about how to detect, measure, and address discriminatory effects that may arise from automated decision-making processes.

Bui et al. (2021) framed the extraction of detailed personal data phrases and associated data collection or sharing practices as a sequence-labeling problem, addressable through an entity-recognition model. The authors developed an entirely automated system named PI-Extract, which uses a neural model to accurately extract privacy practices and significantly outperforms strong rule-based baselines.

Valle-Cruz et al. (2020) aimed at evaluating the public policy-cycle framework in the context of AI, focusing on the actual and anticipated changes that these emerging technologies will introduce at different stages of the policy-making process.

To achieve intelligent analysis of a large number of regulatory policies, Wang et al. (2023) proposes a discourse parsing technique designed for an in-depth understanding of Chinese government documents (CGDs). Utilizing Superstructure Schema and Rhetorical Structure Theory (RST), the paper examines the stylistic characteristics and macrostructure patterns of CGDs, and it develops a discourse analysis framework to define their functional structure and semantic system. Experimental results indicate that the parsing model, which incorporates inherent CGD discourse features, outperforms baseline models. However, despite its high accuracy, the proposed approach may face challenges when applied to cross-format government policies in the real world.

3 Overall Architecture

In this section, we first provide an overview of PolicyInsight's high-level architecture. Then, we dive into the main design decisions in the framework.

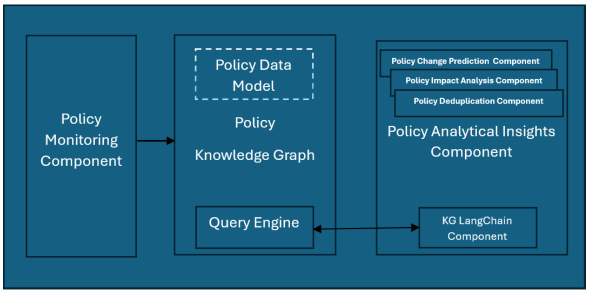

Figure 1: The system architecture of PolicyInsight

3.1 Overview

The PolicyInsight framework is based on four foundational functional components: a dynamic policy data model, a policy knowledge graph built from policy data model entities and relationships, a policy monitoring component and an analytical insights component.

3.2 Policy Data Model

Designing a dynamic data model to represent regulatory policies requires careful consideration of the evolving nature of policies, the diverse range of policy components, and the need for flexibility and scalability. When designing a policy data model, several key considerations must be taken into account to ensure its effectiveness, adaptability, and security. Firstly, it's crucial to identify key entities and attributes within the policy domain, capturing essential elements of policies and their associated metadata. Additionally, defining policy states and incorporating versioning and history tracking mechanisms allows for the monitoring and management of policy changes over time. Finally, prioritizing data integrity and security measures safeguards sensitive policy information, ensuring confidentiality, integrity, and availability throughout the data life cycle.

The policy data model is designed for the bylaws open data (of Ottawa, 2024). A bylaw is a rule or regulation enacted by a local authority, such as a city council or municipal government, to govern conduct, activities, and operations within a specific jurisdiction. Bylaws are subordinate to higher-level laws and are typically enacted to address local issues, maintain order, and regulate various aspects of community life. Moreover, Policies from different government bodies, such as federal, state, and local authorities, are interconnected and often interact in complex ways due to the shared jurisdictional responsibilities, overlapping regulatory

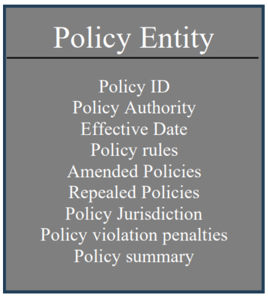

Figure 2: A partial view of Policy Data Model.

Figure 3: Policy lifecycle flow.

frameworks, and intergovernmental relations.

Our policy data model consists of several data entities and their relations (sample in Figure 2). Example entities include Policy entity, Stakeholder entity, Policy Document entity and so on. Similarly, example entity relationships include Policy entity to Stakeholder entity and Policy entity to Policy Document entity. For complete policy data model please refer to project GitHub page.

A typical policy lifecycle consists of several stages or phases that a policy undergoes from its initial conceptualization to its eventual termination or replacement. While the specific stages may vary depending on the context, jurisdiction, and nature of the policy, the following are common stages observed in many policy lifecycles.

3.3 Policy Knowledge Graph

Knowledge graphs in the system help capturing real-time policy data and mitigate issues such as hallucination and poor explainability. Unlike LLMs, which rely on static training data and may generate responses that are not grounded in reality, KGs can be updated in real time to reflect changing policy circumstances. This allows KGs to provide more accurate and reliable information, reducing the risk of hallucination. Additionally, KGs' transparent and interpretable structure enables explainability, as relationships between entities are explicitly defined, making it clear why a particular response was generated. By incorporating real-time data into KGs, organizations can ensure that their decision-making processes are informed by the most up-to-date information, reducing the likelihood of errors and biases associated with LLMs.

PolicyInsight Policy Knowledge Graph is based on the popular graph database Kùzu (Salihoglu, 2023; Inc., 2023), a highly scalable, extremely fast and easy-to-use embeddable database which allows graph-based modeling and querying, graph-optimized storage and graph-optimized query execution. As an extension to the database and querying module, we built a GUI for user input and querying.

Building a knowledge graph in Kùzu from the prepared policy data consists of two primary steps: Creating schema with the designed entities and relationships as Tables and populating tables with prepared CSV data files. As outlined in Section 3.2, which focuses on the design of the policy model schema, we established a data model of entities and their relationships, resulting in the creation of triplets in the form of (entity1, relationship1, entity2) that comprise the knowledge graph. The complete details of Knowledge Graphs schema can be found in the project GitHub page. With the schema fully defined and populated, the knowledge graph is now primed for querying and analysis.

Cypher (Kùzu, 2023) is Kùzu's graph query language that enables data retrieval from the graph. Much like SQL for relational databases, it was inspired by SQL, allowing you to concentrate on the desired data from the graph without worrying about the retrieval process. Given a query objective, like SQL, Cypher also provisions several ways to perform queries to retrieve desired outcome using several languages constructs like query and subquery clauses (Kùzu, 2023).

3.4 Policy Monitoring Component

When designing a policy monitoring component, several critical considerations must be addressed to ensure its effectiveness in tracking policy developments, assessing impacts, and facilitating adaptive governance processes. Firstly, real-time or near real-time updates are essential to provide timely information on policy changes, enabling stakeholders to stay informed and responsive to evolving policy landscapes.

We designed a policy monitoring tool based on a web crawler designed to systematically and automatically collect, analyze, and aggregate policy-related information from various online sources, including regulatory websites, legislative databases, news portals, and other relevant platforms. The tool is configured to identify and prioritize specific sources of policy information, such as regulatory websites, legislative databases, regulatory agencies, and reputable news outlets. This ensures that the collected data is reliable, authoritative, and up to date. Upon extraction, the tool performs content analysis and classification to categorize policy-related information based on predefined topics, keywords, or themes. The tool provides real-time updates and alerts on policy developments, changes, and announcements.

The extracted policy data is stored in a Knowledge Graph for easy access, retrieval, and analysis. Overall, a policy monitoring tool based on a web crawler streamlines the process of collecting, analyzing, and monitoring policy-related information from online sources, empowering policymakers, analysts, and stakeholders to stay informed, responsive, and proactive in addressing policy challenges and opportunities.

3.5 Analytics Insights Component

The primary goal of designing a policy analytical insights component was to enable comprehensive analysis and decision-making support for policymakers and stakeholders. The Analytics Insights Component consists of three subcomponents: Policy Changes Summarization component, Policy Impact Analysis component and Policy Change Predictions component.

Firstly, the component should incorporate policy change summarization capabilities to distill complex policy updates into concise, digestible summaries, facilitating quick understanding of key changes and their implications. Policy changes summarization design flow consists of data preprocessing where the policy documents are preprocessed to remove noise, such as headers, footers, and boilerplate text, and tokenize the text into sentences and paragraphs. And then an LLM was used to generate summaries of policy changes. This involves providing the model with input text (e.g., a section of a policy document) and prompting it to generate a concise summary of the content. The model generates summaries by predicting the most relevant and informative sentences or phrases based on the input context. For the evaluation, the quality of the generated summaries is evaluated using metrics such as ROUGE (Recall-Oriented Understudy for Gisting Evaluation), which measures the overlap between the generated summaries and reference summaries (e.g., human-authored summaries).

Similarly, using an LLM for Policy Impact Analysis involves leveraging its capabilities in NLU and generation to assess the effects and implications of policy interventions. The first step involves gathering relevant data sources, including policy documents, legislative texts, regulatory reports, news articles, and social media discussions related to the policy under analysis. These sources provide context and information about the policy's objectives, implementation, and outcomes. Then, fine-tune a pre-trained LLM on a dataset containing policy-related texts and documents. Provide prompts or queries to the fine-tuned LLM to prompt it to generate assessments or predictions about the policy's impact. For example, prompt the model with questions such as "What are the potential economic effects of implementing this policy?" or "How might this policy impact different demographic groups?" The LLM generates impact analyses by predicting potential outcomes, consequences, and implications of the policy under consideration. Evaluate the quality and validity of the generated impact analyses using expert review, validation against empirical data, or comparison with existing impact assessments.

For Policy Change Predictions, the designed workflow involves gathering a comprehensive dataset of historical policy documents, legislative texts, regulatory reports, news articles, and other relevant sources that document past policy changes and developments. This dataset serves as the training data for the LLM. Then, fine-tune a pre-trained LLM on the historical policy dataset. For example, prompt the model with questions such as "What policy changes are likely to occur in the next year based on historical trends?" or "Which policy areas are expected to see significant changes?" The LLM generates policy change predictions by analyzing patterns, correlations, and signals in the historical data. Evaluate the quality and accuracy of the generated policy change predictions using metrics such as precision, recall, and F1-score. Validate the predictions against empirical data or expert judgments to assess their reliability and usefulness.

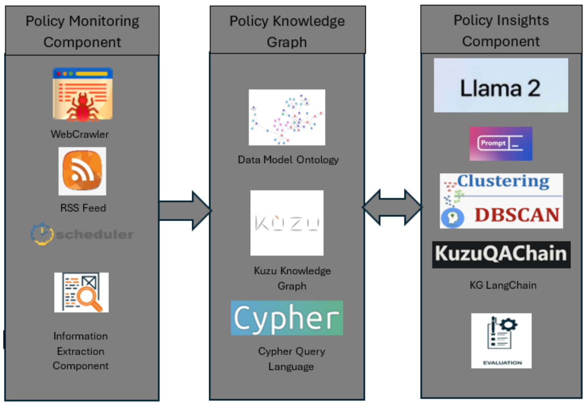

Figure 4: An overview of PolicyInsight.

4 Implementation Paradigm

The overall structure of PolicyInsight and the connections between different modules are illustrated in Figure 4. In this section, we discuss the detailed implementation of three main components of PolicyInsight: the monitoring subsystem, streaming data management subsystem, and the three-layered monitoring subsystem. We show how to combine different technologies to achieve a high-performance data-analytics system for PolicyInsight.

4.1 Policy Monitoring Component

At the core of the system is a custom web crawler designed to efficiently traverse regulatory websites, regulatory portals, legislative databases, and other online sources to collect policy-related data. The web crawler employs intelligent algorithms to navigate complex website structures, extract relevant information, and filter out noise and irrelevant content. Depth-First crawling strategy was used with Time-based rate limiting considering the overnight update of policies.

In addition to web crawling, the system incorporates RSS feed mechanisms to subscribe to policy-related feeds from authoritative sources, government agencies, industry publications, and news outlets.

Additionally, integrating Llama 2, an LLM, with suitable prompts further enhances the system's capabilities. Llama 2 can be utilized for NLP tasks such as policy summary generation, obligations detection, and risks identification. Leveraging LLM's capabilities allows for comprehensive analysis of policy text, enabling the generation of concise summaries and the extraction of obligations (e.g., regulatory requirements, compliance mandates) and potential risks associated with policy provisions.

To maintain data integrity and reliability, quality assurance measures are implemented to validate the accuracy, completeness, and relevance of extracted policy insights. Validation checks, error handling mechanisms, and human-in-the-loop review processes are incorporated to ensure the reliability and integrity of the output generated by the system.

4.2 Policy Knowledge Graph

In the process of building a policy knowledge graph, the system leverages a pre-designed policy data model to structure the information extracted from the JSON output generated in the previous step of the policy monitoring component. This pre-designed data model serves as a blueprint for organizing policy-related entities, relationships, and attributes in a structured and consistent manner.

The first crucial step in this process involves mapping the entities identified in the JSON output to the corresponding entity types defined in the policy data model. Entities such as policies, regulations, stakeholders, and risks are matched with their counterparts in the data model, ensuring alignment between the extracted information and the predefined entity schema.

Once the entities are mapped, the system proceeds to establish relationships between them based on the predefined relationship types defined in the policy data model. Relationships such as "is_related_to", "imposes_obligation_on," and "addresses_risk" are identified and established between entities, capturing the connections and dependencies between different policy elements.

With the entities and relationships mapped and established, the system populates the knowledge graph, accordingly, creating nodes for each entity type and edges for each relationship type.

In implementing Cypher queries to extract crucial insights from the policy knowledge graph, the system capitalizes on the expressive capabilities of Cypher, a graph query language specifically designed for graph databases.

4.3 Policy Analytical Insights Component

The implementation of the policy analytical insights component leverages the emergent abilities of LLMs to analyze extensive repositories of policy documents, legislative texts, and regulatory frameworks.

4.3.1 Implementing Policy Change Summarization Component

The implementation of policy change summarization began with the extraction of article summaries using LLM Llama 2 (Touvron et al., 2023) by applying prompt techniques (Liu et al., 2023) (Varadarajan and Hristidis, 2006), which enabled the system to distill key insights and highlights from a vast array of policy documents and legislative texts. However, the initial approach of clustering these summaries led to a significant number of false positives, as similar policy articles were erroneously grouped together due to semantic overlaps or contextual similarities. To address this challenge, the system augmented the policy summary data with rich metadata sourced from the knowledge graph, a technique known as KG-enhanced LLMs (Pan et al., 2024), encompassing attributes such as policy maker, jurisdiction, regulatory domain, and effective date.

To incorporate policy metadata from the knowledge graph into the summaries generated by LLM we used a LangChain (Topsakal and Akinci, 2023) based tool for KuzuDB called KuzuQAChain (langchain ai, 2024), so that the system can gain additional contextual information and domain-specific insights that facilitated more accurate deduplication of policy articles.

Through this iterative approach, the system achieved a significant reduction in false positives and improved the accuracy of policy deduplication by leveraging the complementary capabilities of article summaries and policy metadata from the knowledge graph.

4.3.2 Implementing Policy Impact Analysis Component

Policy impact analysis was implemented through a multi-faceted approach that began with the generation of policy core areas or topics derived from the analysis of policy documents and regulatory frameworks using Llama 2 using appropriate prompt technique (Liu et al., 2023) (Varadarajan and Hristidis, 2006). Subsequently, following a different approach of LLM-augmented KGs to unify LLMs with KGs (Pan et al., 2024), these policy core areas were stored within a knowledge graph, enriching the graph with contextual information and semantic relationships that facilitated comprehensive impact analysis.

4.3.3 Implementing Policy Prediction Component

The policy prediction component was implemented to harness the synergistic capabilities of both LLMs and knowledge graphs by using a technique called Synergized LLMs + KGs (Pan et al., 2024), for the predictive analytics in the policy domain. At its core, this component employed advanced NLP techniques powered by Llama 2 to analyze vast repositories of unstructured textual data comprising policy documents, legislative texts, and regulatory frameworks. By training on historical policy data and learning from nuanced linguistic patterns, Llama 2 could generate plausible scenarios, anticipate emerging policy trends, and forecast future regulatory changes with remarkable accuracy.

4.4 Evaluation Results

In this section, we would like to present evaluation results of two use cases to assess the efficacy of unifying the capabilities of LLMs and Knowledge Graphs in policy analysis which revealed remarkably high accuracy results for both the policy deduplication and policy impact analysis tasks.

Use case 1: Policy deduplication results Objective: The primary objective of this task is to identify and remove duplicate policies from a dataset containing policies from overlapping jurisdictions but serving the same purpose.

Test Data: The test data comprises a curated selection of policy samples sourced from overlapping jurisdictions, enacted for both similar and disparate purposes, and meticulously hand-labeled for evaluation purposes. Experiment 1 approach: DBScan clustering was performed to cluster policy summaries generated by employing Llams 2 with prompts.

Results: An overall accuracy of 85% was achieved by using LLMs only due to the huge number of false positives (Figure 5).

Experiment 2 approach: DBScan clustering was performed to cluster policy summaries generated by employing Llams 2 with prompts. But this time policy summaries are augmented with corresponding policy metadata like policy maker and jurisdiction, policy effective data etc. coming from the policy knowledge graph.

Results: Overall accuracy was boosted to 95% with a reduced number of false positives (Figure 6).

Use case 2: Policy Impact Analysis results

Objective: The primary objective of this task

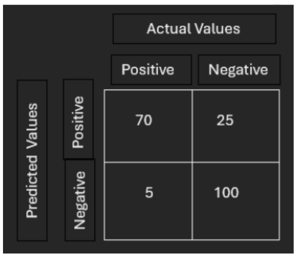

Figure 5: Confusion matrix when only LLM capabilities are employed for policy deduplication.

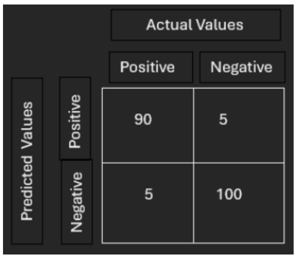

Figure 6: Confusion matrix when KG-enhanced LLMs are employed for policy deduplication.

is to identify the customers impacted by a policy change.

Test Data: The test data comprises a curated selection of policy samples sourced from a policy body and labeled automatically using breadcrumb approach for evaluation purposes.

Experiment approach: LLM-augmented KGs approach was employed in which Llama 2 was used with prompt to identify key impacted areas of a given policy and fed that information to Policy KG along with other derived policy information. During inference, Policy KG was queried to match with customer business domains to identify impacted customers.

Results: An overall accuracy of 89% was achieved by this approach.

5 Conclusions

In conclusion, our work introduces PolicyInsight, a novel analytical framework designed to address the evolving challenges of regulatory document monitoring in a rapidly changing socio-economic landscape. By leveraging a sophisticated policy data model and state-of-the-art NLP and knowledge graph techniques in a combined fashion, PolicyInsight enables stakeholders to continuous monitoring and derive timely insights from policy documents, fostering data-driven decision-making. Incorporating a novel dynamic policy data model for a scalable and efficient knowledge graph, PolicyInsight leverages an innovative unified approach to combining capabilities of both LLMs and KGs to achieve remarkable accuracy for policy deduplication, policy impact analysis and policy changes prediction. By providing stakeholders with access to actionable insights derived from policy data, PolicyInsight empowers policymakers, businesses, and citizens to make informed decisions, respond effectively to crises, and comply with regulatory requirements. Looking ahead, the continued refinement and expansion of PolicyInsight holds immense potential for driving positive change in governance practices. Future research endeavors may focus on enhancing the scalability, interoperability, and predictive capabilities of PolicyInsight, thereby enabling stakeholders to anticipate regulatory changes, identify emerging trends, and proactively address societal challenges. Our future work also addresses the few remaining items from the framework. In summary, PolicyInsight stands at the forefront of innovation in policy monitoring and analysis, offering a powerful tool for navigating the complexities of the modern regulatory landscape and fostering a more transparent, informed, and responsive governance ecosystem.

References

Bilal Abu-Salih. 2021. Domain-specific knowledge graphs: A survey. *Journal of Network and Computer Applications*, 185:103076.

Duc Viet Bui, Kang G. Shin, Jong-Min Choi, and Junbum Shin. 2021. Automated extraction and presentation of data practices in privacy policies. *Proceedings on Privacy Enhancing Technologies*, 2021:88 – 110.

Danilo Dessì, Francesco Osborne, Diego Reforgiato Recupero, Davide Buscaldi, and Enrico Motta. 2021. Generating knowledge graphs by employing natural language processing and machine learning techniques within the scholarly domain. *Future Generation Computer Systems*, 116:253–264.

V. Devedžić. 1999. A survey of modern knowledge modeling techniques. *Expert Systems with Applications*, 17(4):275–294.

Kùzu Inc. 2023. Kùzu: A highly scalable, extremely fast, easy-to-use embeddable graph database. https://kuzudb.com/.

Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and Philip S. Yu. 2022. A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2):494–514.

Peter Parycek Johann Höchtl and Ralph Schöllhammer. 2016. Big data in the policy cycle: Policy decision making in the digital era. *Journal of Organizational Computing and Electronic Commerce*, 26(1-2):147–169.

Cypher: Kùzu. 2023. Cypher manual. https://docs.kuzudb.com/cypher/.

langchain ai. 2024. Langchain kùzu. https://python.langchain.com/docs/integrations/graphs/kuzu_db/.

Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Comput. Surv.*, 55(9).

City of Ottawa. 2024. By-laws, licences and permits. https://ottawa.ca/en/living-ottawa/laws-licences-and-permits.

Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Jiapu Wang, and Xindong Wu. 2024. Unifying large language models and knowledge graphs: A roadmap. *IEEE Transactions on Knowledge and Data Engineering*, 36(7):3580–3599.

Samira Pouyanfar, Saad Sadiq, Yilin Yan, Haiman Tian, Yudong Tao, Maria Presa Reyes, Mei-Ling Shyu, Shu-Ching Chen, and S. S. Iyengar. 2018. A survey on deep learning: Algorithms, techniques, and applications. *ACM Comput. Surv.*, 51(5).

Semih Salihoglu. 2023. Kùzu: A database management system for "beyond relational" workloads. *SIGMOD Rec.*, 52(3):39–40.

Phillip Schneider, Tim Schopf, Juraj Vladika, Mikhail Galkin, Elena Simperl, and Florian Matthes. 2022. A decade of knowledge graphs in natural language processing: A survey. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 601–614, Online only. Association for Computational Linguistics.

Oguzhan Topsakal and T. Cetin Akinci. 2023. Creating large language model applications utilizing langchain: A primer on developing llm apps fast. *International Conference on Applied Engineering and Natural Sciences*, 1:1050–1056.

Hugo Touvron, Louis Martin, Kevin R. Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Daniel M. Bikel, Lukas Blecher, Cristian Cantón Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony S. Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel M. Kloumann, A. V. Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, R. Subramanian, Xia Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zhengxu Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melissa Hall Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open foundation and fine-tuned chat models. *ArXiv*, abs/2307.09288.

David Valle-Cruz, J. Ignacio Criado, Rodrigo Sandoval-Almazán, and Edgar A. Ruvalcaba-Gomez. 2020. Assessing the public policy-cycle framework in the age of artificial intelligence: From agenda-setting to policy evaluation. *Government Information Quarterly*, 37(4):101509.

Ramakrishna Varadarajan and Vagelis Hristidis. 2006. A system for query-specific document summarization. In *Proceedings of the 15th ACM International Conference on Information and Knowledge Management*, CIKM '06, page 622–631, New York, NY, USA. Association for Computing Machinery.

Xiaoyu Wang, Hong Zhao, Hongzhi Zhu, and Fang Wang. 2023. Towards intelligent policy analysis: A discourse structure parsing technique for chinese government document. *Information Processing & Management*, 60(4):103363.

Richard W Waterman and B Dan Wood. 1993. Policy monitoring and policy analysis. *Journal of Policy Analysis and Management*, 12(4):685–699.

Ming Zhou, Nan Duan, Shujie Liu, and Heung-Yeung Shum. 2020. Progress in neural nlp: Modeling, learning, and reasoning. *Engineering*, 6(3):275–290.