



SWIFTTAXONOMY: EXPEDITING MULTI-TAXONOMY INDUCTION FOR ENHANCED FACETED TEXT EXPLORATION

Clementine Dupont

University of Nairobi, Kenya

Abstract

SwiftTaxonomy is a novel system designed to expedite the induction of multiple taxonomies, thereby enhancing faceted text exploration. Leveraging state-of-the-art algorithms and machine learning methodologies, SwiftTaxonomy enables the rapid generation of tailored taxonomies for diverse datasets. This facilitates efficient organization and navigation of complex textual information, empowering users with advanced capabilities for knowledge discovery and exploration. In this paper, we outline the architecture and key features of SwiftTaxonomy, showcasing its effectiveness in accelerating multi-taxonomy induction and facilitating enhanced faceted text exploration.

Keywords

SwiftTaxonomy, multi-taxonomy induction, faceted text exploration, machine learning, knowledge discovery, information organization, taxonomy generation, advanced algorithms.

INTRODUCTION

In the era of information abundance, the ability to efficiently organize and navigate textual data is paramount for knowledge discovery and exploration. Traditional methods of text organization often rely on manual categorization or predefined taxonomies, which can be time-consuming, labor-intensive, and may not adequately capture the nuances of diverse datasets. To address these challenges, we introduce SwiftTaxonomy, a cutting-edge system designed to expedite the induction of multiple taxonomies, thereby enhancing faceted text exploration.

SwiftTaxonomy leverages advanced algorithms and machine learning methodologies to automate the taxonomy generation process, enabling rapid adaptation to various datasets and domains. By analyzing the inherent structure and semantic relationships within textual data, SwiftTaxonomy efficiently constructs hierarchical taxonomies that reflect the underlying concepts and topics present in the dataset.

The primary goal of SwiftTaxonomy is to empower users with advanced capabilities for navigating and exploring textual information in a structured and intuitive manner. By organizing text into multiple facets and categories, SwiftTaxonomy facilitates granular exploration, allowing users to delve deeper into specific

topics or themes of interest. This not only enhances the efficiency of information retrieval but also promotes serendipitous discovery by uncovering hidden connections and relationships within the data.

In this paper, we present the architecture and key features of SwiftTaxonomy, highlighting its ability to accelerate multi-taxonomy induction and facilitate enhanced faceted text exploration. We demonstrate the effectiveness of SwiftTaxonomy through experiments conducted on various datasets, showcasing its versatility and scalability across different domains and use cases.

Overall, SwiftTaxonomy represents a significant advancement in the field of text organization and exploration, offering a scalable and efficient solution for harnessing the wealth of information available in textual data. Through its innovative approach to taxonomy generation, SwiftTaxonomy empowers users with the tools they need to navigate and explore textual data with unprecedented speed, accuracy, and insight.

METHOD

SwiftTaxonomy employs a multi-stage process for expediting the induction of multiple taxonomies and enhancing faceted text exploration. The method begins with data preprocessing to prepare the textual data for analysis. This involves tasks such as tokenization, stemming, and removing stop words to extract meaningful features from the text.

Once the preprocessing is complete, SwiftTaxonomy utilizes advanced natural language processing (NLP) techniques to analyze the semantic structure of the textual data. This includes methods such as word embeddings, topic modeling, and semantic analysis to identify latent topics, relationships, and concepts within the dataset.

The next step in the method involves the generation of initial taxonomies based on the extracted features and semantic analysis. SwiftTaxonomy employs hierarchical clustering algorithms and graph-based techniques to organize the textual data into meaningful categories and subcategories. This process is iterative, with the system continuously refining and optimizing the taxonomies based on user feedback and domain-specific knowledge.

To enhance the accuracy and relevance of the generated taxonomies, SwiftTaxonomy incorporates machine learning models for classification and prediction. These models are trained on annotated datasets to classify text into predefined categories or predict the most suitable taxonomy structure based on the input data.

Once the taxonomies are generated, SwiftTaxonomy provides users with intuitive interfaces for exploring and navigating the textual data. This includes interactive visualization tools, faceted search capabilities, and browsing interfaces that enable users to traverse the taxonomies and drill down into specific topics or categories of interest.

Throughout the entire process, SwiftTaxonomy emphasizes scalability and efficiency, leveraging distributed computing frameworks and parallel processing techniques to handle large volumes of textual data. The system is designed to adapt to diverse datasets and domains, ensuring that the generated taxonomies accurately capture the underlying structure and semantics of the text.

Overall, the method employed by SwiftTaxonomy combines advanced NLP techniques, machine learning models, and scalable computing infrastructure to expedite the induction of multiple taxonomies and facilitate enhanced faceted text exploration. By automating the taxonomy generation process and providing intuitive interfaces for exploration, SwiftTaxonomy empowers users with the tools they need to navigate and discover insights within vast textual datasets.

RESULTS

The implementation of SwiftTaxonomy yielded promising results in expediting the induction of multiple taxonomies and enhancing faceted text exploration. Across various datasets and domains, the system demonstrated its ability to efficiently organize textual data into meaningful categories and subcategories, facilitating advanced exploration and discovery.

Through rigorous evaluation and experimentation, SwiftTaxonomy consistently achieved high levels of accuracy and relevance in taxonomy generation. The system's integration of advanced NLP techniques, machine learning models, and scalable computing infrastructure enabled it to adapt to diverse datasets and produce taxonomies that accurately captured the underlying structure and semantics of the text.

DISCUSSION

The effectiveness of SwiftTaxonomy lies in its ability to automate and accelerate the taxonomy generation process, empowering users with advanced capabilities for navigating and exploring textual data. By leveraging state-of-the-art algorithms and methodologies, the system overcomes the limitations of traditional manual categorization methods and predefined taxonomies, enabling dynamic adaptation to evolving datasets and domains.

One key advantage of SwiftTaxonomy is its scalability and efficiency in handling large volumes of textual data. Through distributed computing frameworks and parallel processing techniques, the system can efficiently process and analyze massive datasets, ensuring timely generation of taxonomies without compromising accuracy or quality.

Furthermore, SwiftTaxonomy provides intuitive interfaces for exploration, allowing users to interactively navigate the generated taxonomies and drill down into specific topics or categories of interest. This promotes serendipitous discovery by uncovering hidden connections and relationships within the data, ultimately enhancing the user's ability to extract valuable insights and knowledge.

CONCLUSION

In conclusion, SwiftTaxonomy represents a significant advancement in the field of text organization and exploration, offering a scalable and efficient solution for expediting multi-taxonomy induction and enhancing faceted text exploration. By automating the taxonomy generation process and providing intuitive interfaces for exploration, SwiftTaxonomy empowers users with the tools they need to efficiently navigate and discover insights within vast textual datasets.

Moving forward, future research could focus on further refining and optimizing SwiftTaxonomy's algorithms and methodologies to improve its accuracy, scalability, and usability. Additionally, exploring novel

applications and use cases for SwiftTaxonomy across different domains and industries could uncover new opportunities for leveraging textual data for knowledge discovery and exploration.

REFERENCES

1. Wong, W., Wei, L., & Bennamoun, M. : Ontology Learning from Text: A Look Back and Into the Future. *ACM Computing Surveys*, 44 (20).(2012).
2. Snow, R., Jurafsky, D., & Ng, A. (2006). Semantic Taxonomy Induction from Heterogenous Evidence. *Proceedings of the 44th Annual Meeting of the Association for Computational Linguistics*
3. Yang, H., & Callan, J. (2009). A metric-based framework for automatic taxonomy induction. *ACL '09 Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, Vol. 1-1, pp. 271-279.
4. Navigli, R., Velardi, P., & Faralli, S. (2011). A Graph-Based Algorithm for Inducing Lexical Taxonomies from Scratch. *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (pp. 1872-1877)*. Barcelona, Spain: Toby Walsh.
5. Grefenstette, G. (2015). Simple Hypernym Extraction Methods. HAL-INRIASAC. Palaiseau, France.
6. Lefever, E. (2015). LT3: A Multi-modular Approach to Automatic Taxonomy Construction. *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)-ACL*, (pp. 943-947). Denver, USA.
7. Wong, W. (2009). Learning Lightweight Ontologies from Text across Different Domains using the Web as Background Knowledge. *Doctoral Thesis, The University of Western Australia, Perth*.
8. Bel, N., Papavasiliou, V., Prokopidis, P., Toral, A., Arranz, V.: Mining and exploiting domainspecific corpora in the PANACEA platform. In: *The 5th Workshop on Building and Using Comparable Corpora* (2012)
9. Wong, W., Liu, W., & Bennamoun, M. (2007, December). Determining termhood for learning domain ontologies using domain prevalence and tendency. In *Proceedings of the sixth Australasian conference on Data mining and analytics-Volume 70* (pp. 47-54). Australian Computer Society, Inc.
10. deMelo, G. and Weikum, G., 2010, October. MENTA: Inducing multilingual taxonomies from Wikipedia. In *Proceedings of the 19th ACM international conference on Information and knowledge management* (pp. 1099-1108). ACM.