CBIE 2023
CONGRESSO BRASILEIRO DE INFORMÁTICA NA EDUCAÇÃO
Uma escola para o futuro: Tecnologia e conectividade a serviço da educação

From Text to Maps: Automated Concept Map Generation Using Fine-tuned Large Language Model

Wagner de A. Perin[1], **Davidsom Cury**[1], **Camila Zacché de Aguiar**[1], **Crediné S. de Menezes**[2]

[1]Departamento de Informática – Universidade Federal do Espírito Santo (UFES)
CT-7 – Av. Fernando Ferrari, 514 – 29075-910 – Vitória – ES – Brazil

[2]Faculdade de Educação – Universidade Federal do Rio Grande do Sul (UFRGS)
Av. Paulo Gama, s/nº – 90040-060 – Porto Alegre – RS – Brazil

{wagner.perin, davidsom.cury, camila.z.aguiar, credine.menezes}@ufes.br

Abstract. *Concept maps (CMs) are tools for visualizing relationships between ideas, facilitating more effective comprehension and learning. However, the automatic generation of CMs from unstructured text presents a challenge, often requiring semantic markup and subsequent complex processing. This paper introduces a novel approach to address this hurdle by harnessing the capabilities of fine-tuned Large Language Models (LLMs). Our innovative methodology uses these models to extract structured propositions from unstructured text, subsequently serving as the foundation for constructing a CM. This process reverses the transformation of CM relations into first-order logic propositions, a concept explored in our previous work. To achieve this, we train the LLM using fine-tuning techniques, leveraging the latest advancements in artificial intelligence and machine learning. We evaluate our proposed solution based on precision and recall metrics, comparing our outcomes against models crafted by experts. Notably, the results indicate that our method can contribute significantly to advancements in the automatic generation of CMs, illustrating another application bolstered by recent breakthroughs in artificial intelligence. As a stepping stone in this promising direction, future research should continue to refine the model and explore potential applications across diverse domains.*

1. Introduction

Concept maps (CMs) are graphical tools for organizing and representing knowledge. They play a key role in the more profound understanding of complex information, allowing students and researchers to visually represent the relationships and connections between concepts, helping them understand, communicate and retain knowledge. However, manually creating CMs from unstructured text is time-consuming and cognitively demanding, requiring substantial experience and effort.

Recent years have demonstrated the incredible potential of artificial intelligence (AI), primarily related to Natural Language Processing (NLP), with considerable advances, particularly with the advent of Large Language Models (LLMs). These models increasingly demonstrate an unprecedented ability to understand and generate human-like text, opening space for many applications, from chatbots and content generation to more sophisticated tasks, such as summarizing and translating text.

Despite these advances, the automatic generation of CMs from unstructured text remains a significant challenge. Unsurprisingly, a secondary study by [Aguiar et al. 2016] identified several techniques and proposed one to assist in automatically constructing CMs from unstructured texts. These processes usually involve semantic markings, followed by complex processing to identify and extract the elements present in the text. However, all these studies realized that the vast variability and ambiguity inherent in natural language make it particularly challenging to accurately extract the meaningful relationships between concepts that form the basis of a CM.

In light of these challenges, this study explores a new approach to automate the creation of CMs from textual data, leveraging the features of fine-tuned LLMs. Specifically, we intend to leverage these models to extract structured propositions from unstructured text, thus providing the necessary fundamental elements for constructing a CM. This strategy will reverse the transformation of CMs into first-order logical propositions, concepts explored in our previous research [Perin et al. 2015].

We used precision and recall metrics, comparing the results with models created by specialists and results reported in [Aguiar et al. 2016]. From the results presented in section 4, our approach has demonstrated a significant improvement over the previously established methods by utilizing a fine-tuned LLM. Regarding concept identification, our approach outperformed by approximately 23.5% in precision and held comparable recall values. Even more notably, in proposition identification, we saw an increase of 121% in precision and a 33.2% improvement in recall compared to the previous methods. These numbers highlight a substantial increase in the effectiveness of our methodology, showcasing the transformative potential of applying LLMs in the realm of automatic CM generation. Ultimately, this work significantly contributes to the automatic CM generation field, highlighting the potential of recent advances in AI to boost this field.

This work will be developed as follows: Section 2 will review the background concepts and related work in the field; Section 3 will detail our proposed methodology and its implementation; Section 4 will discuss the results and their implications respectively; and, finally, Section 5 will conclude the article and propose directions for future research.

2. Related Works and Background

The automatic construction of CMs from unstructured texts has been a research topic of great interest for many years. Based on a comprehensive literature review conducted between 1994 and 2016, various technological approaches for automatic CM construction were identified and categorized [Aguiar et al. 2016]. Four notable works stood out among these approaches, each offering unique methods for generating CMs [Wang et al. 2008, Žubrinić et al. 2015, Zouaq and Nkambou 2009, de la Villa et al. 2012]. The underlying methodology used across these approaches typically involved morphological and syntactic analysis, semantic mapping, statistical analysis, and linguistic techniques. However, they highlighted several challenges: sentence fragmentation, long or missing concept labels, and dependence on other data sources like ontologies, knowledge bases, or thesauri.

In parallel, data mining methods have been used effectively to unearth useful information in large datasets across various fields [Booth 2007, Li et al. 2004, Cowan 2002, Hirschman et al. 2002]. Despite its later adoption in education, educational data mining has shown great promise in designing smarter learning technologies

[Romero and Ventura 2007, Baker 2014]. In this context, CMs have been leveraged to represent knowledge structures, with several research efforts being directed towards generating these maps [Acharya and Sinha 2017, Tseng et al. 2007, Bai and Chen 2008, Chen and Sue 2013, Oppl and Stary 2011, Acharya and Sinha 2017].

Despite the various algorithms and promising results achieved by these studies, they often relied on explicit final datasets. They also rarely delved into the challenges of handling multi-semantic data like raw textual data, which takes time to acquire. In contrast, many studies have focused on text analysis and automatic CM generation [Lai et al. 2017, Wang and Hu 2016, Qasim et al. 2013, Nugumanova et al. 2015]. These studies have demonstrated the ability to save time by effectively dealing with unstructured texts, thus expediting the process of CM generation [Shao et al. 2020].

This article presents a new approach for automatic CM construction using fine-tuned LLMs to face some of the limitations identified in our previous works on NLP and the specific challenges of this domain, taking advantage of the opportunities offered by the recent advances in AI technologies, especially concerning NLP[Perin et al. 2014].

2.1. Fine-Tuning of Large Language Models (LLMs)

Large Language Models(LLMs) are AI models trained on massive amounts of data to understand and generate human-like text. They can perform tasks like language translation, question answering, and summarization. However, while these pre-trained LLMs are powerful, they may need to be optimized for specific applications or domains. In such instances, fine-tuning is employed to enhance the LLM's performance.

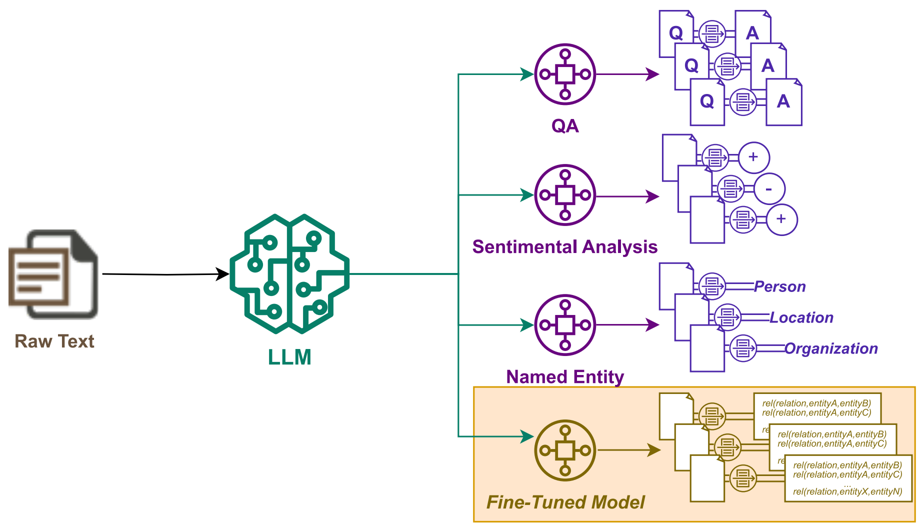

Figura 1. Fine-tuning of a pre-existing LLM for specific purposes

As illustrated in Fig.1, fine-tuning involves retraining a pre-existing model on new, task-specific data, enhancing its capability to handle specialized tasks or domains. This process includes loading the pre-trained model, adjusting its upper layers for the new task, and freezing the lower layers to maintain the general features already learned.

Upon training the new layers, the whole model undergoes fine-tuning using task-specific data. This process customizes the model to the new task, improving its performance and accuracy, thereby allowing it to serve domain-specific applications effectively.

The following section will detail the proposed architecture for leveraging fine-tuned LLMs in automatically generating CMs from unstructured text.

3. Methodology and Implementation

LLMs demonstrate remarkable capabilities in processing and interpreting complex language structures, recognizing sentence structures, identifying named entities, syntactic structures, and much more. Our methodology leverages these capabilities to understand unstructured text, identify pertinent propositions, and produce structured data that subsequently form the basis of a CM. This section elucidates how these fine-tuned LLMs can replace complex processing algorithms and streamline the transformation from unstructured texts to structured ones.

3.1. Conceptual Architecture: A Comparative Overview

A proposition comprises a tuple (concept-relation-concept), where the relational term (often a verb) outlines the semantic link between two concepts. Some studies use the propositions in CMs as a knowledge base in intelligent systems [Perin and Cury 2016].

As mentioned, the automatic generation of CMs from unstructured texts requires a sophisticated network of specific algorithms to detect and treat textual elements. The central principle that drives these algorithms is identifying and extracting these tuples, or propositions, forming the fundamental structures for constructing the corresponding CM. Therefore, an overview of automatic CM generators will invariably reveal components analogous to those depicted in Fig.2(a).

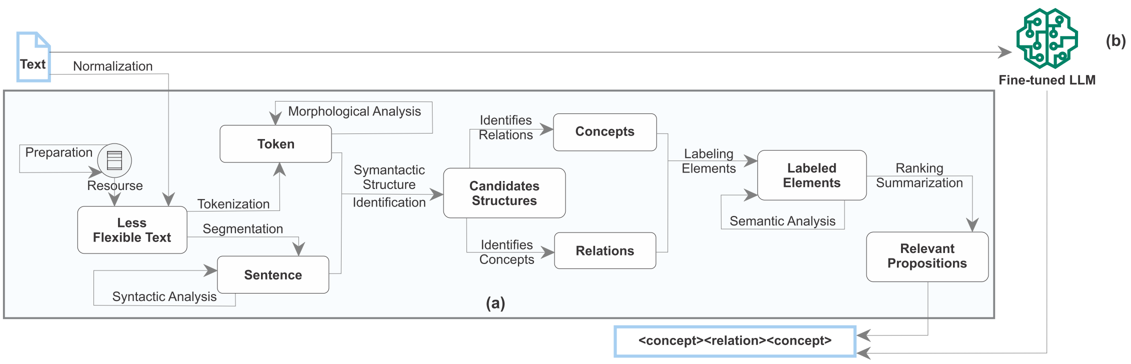

Figura 2. (a) Classical Architecture of CM Generators [Aguiar et al. 2016]. (b) Proposed Architecture.

In Fig.2(a), procedures inscribed in the arrows represent specialized algorithms designed for the text's treatment, annotation, identification, and interpretation, to reach the ultimate goal of extracting the propositions that will constitute the generated CM. In addition to being computationally complex, these algorithms depend on auxiliary databases: in that case, the authors use WordNet[1] and DBpedia[2], which, together with SPARQL[3], support the process of 'labeling' and 'semantic analysis.'

---

[1] **WordNet** is an extensive lexical database of English. Nouns, verbs, adjectives, and adverbs are grouped into sets of cognitive synonyms (synsets), each expressing a distinct concept.

[2] **DBpedia** is a project that aims to extract structured content from the information created as part of the Wikipedia project.

[3] **SPARQL** is a query language and protocol for semantic web data sources used to retrieve information from databases stored in Resource Description Framework (RDF) format.

In classical architecture, every specialized algorithm has a specific role in processing text. However, as a set, they work towards one goal: converting raw input text into a set of extracted propositions.

Considering that LLMs are inherently equipped with the necessary 'capabilities' for this task, the proposed architecture—depicted in Fig.2(b)—replaces these complex algorithms with a fine-tuned LLM. This template's goal is straightforward: direct raw text input into a list of propositions. This means that the complexities associated with managing complex structures and variations in syntax and semantics in unstructured text are handled by the lower layers of LLM. The upper layers, in turn, are tuned to extract propositions specifically (see Fig.1).

3.2. Implementation

To implement the proposed architecture, we followed a four-step process whose interrelationships can be seen in Fig.3. Each process plays a crucial role in the overall architectural picture. In the following subsections, we present some important details of each process, its sub-processes, inputs, and artifacts:

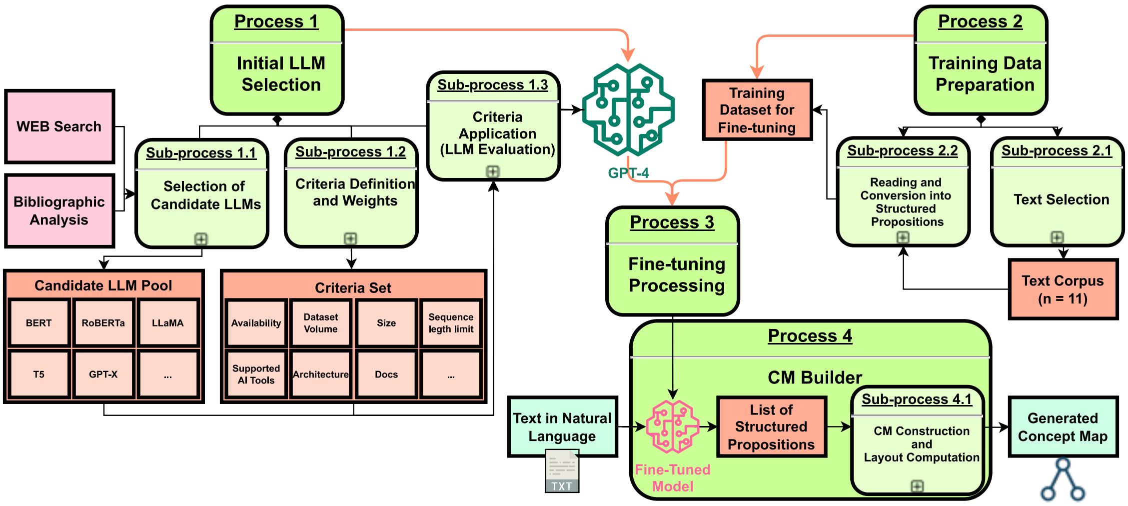

Figura 3. Conceptual Architecture Overview

3.2.1. Process 1: Initial LLM Selection

LLM selection is critical as it sets the stage for the rest of the implementation process. In our proposed architecture, this first process revolves around analyzing and choosing the LLM as the foundation for the subsequent fine-tuning stage (Process 3). This process was subdivided into three distinct sub-processes:

- **Selection of Candidate LLMs**: entailed reviewing a comprehensive secondary study[Yang et al. 2023] and extensive web-based searches to identify existing LLMs readily accessible online. As an artifact of this sub-process, we generate a tabulated inventory called the *LLM Candidate Pool (n=52)*.
- **Criteria Definition and Weights**: we define the crucial criteria list to select the most appropriate LLM for our project. Each criterion was assigned a weight ranging between 0 and 5, denoting its significance to our objectives. Consequently,

this sub-process yielded an artifact called the *Criteria Set*, which comprises a weighted set of criteria for the next sub-process (LLM evaluation).

- **Criteria Application (LLM Evaluation**: was conducted in two steps. ***Step 1***: the most important criteria (4 or 5 in the *Criteria Set*) were used to filter the *Candidate LLM Pool*, reducing it from 52 to 9. ***Step 2***: These latter were evaluated with the remaining criteria through a scoring system, producing a final ranking.

Following this analysis, we settled on GPT-4 as the LLM of choice for our project [Ecoffet 2023]. This decision has been mainly influenced by its superior understanding and generation of coherent text, capacity to identify named entities, and capability to interpret complex sentences and syntax: crucial attributes for automatic proposition extraction from unstructured text. Also, its popularity, extensive documentation, and access to pre-trained models and fine-tuning tutorials further facilitated our implementation. Finally, the vast training dataset and high grammatical and factual accuracy made it the most suitable LLM for our project's requirements.

3.2.2. Process 2: Training Data Preparation

This process involves creating a dataset with examples of expected inputs and outputs. This dataset lays the groundwork for the subsequent fine-tuning process, as the model learns to perform the specific task by observing the patterns in this dataset. We divide this process into two sub-processes, namely:

- **Text Selection**: we chose a range of texts to form the basis for the dataset construction. Despite GPT-4's capability of handling inputs up to 25,000 words, our team selected 11 texts with a limit of 4,000 words each. This decision was based on the size of our research team and the subsequent sub-process, which demands substantial cognitive effort from the research team. Also, higher-capacity models, such as GPT-4, are trained on large volumes of data with multiple layers, requiring fewer examples in fine-tuning, facilitating the dataset construction process.
- **Reading and Conversion into Structured Propositions**: our researchers meticulously read the selected texts, identify the inherent propositions, and transform these into a structured format akin to the approach adopted in our prior research [Perin 2014]. Consequently, from the 11 chosen texts, 387 propositions were discerned. So, each text was individually associated with its corresponding set of propositions, forming the *Training Dataset*.

3.2.3. Process 3: Fine-tuning Process

This process focuses on customizing the LLM to accomplish a specific task, made possible by leveraging the dataset created in the previous process as a guide for further refinement of the LLM. The training process involves feeding this dataset into the model, allowing the machine learning engine to identify and adapt to the patterns and structures characteristic of our task. The result is a new variant of the LLM, fine-tuned and uniquely suited for the task: extracting propositions from unstructured texts. This process is critical, as it tailors the broad capabilities of the general LLM to our specific requirement, enhancing its performance and efficiency in generating structured data for CM generation.

3.2.4. Process 4: CM Builder

This process utilizes the *Fine-Tuned Model* as an initial component, responsible for performing the entire NLP task: primarily extracting a list of propositions from the input unstructured text. This list represents the key concepts and their relationships, essentially the 'knowledge' contained within the text. They form the backbone of our CM, informing the identification and extraction of elements to be visualized as nodes (concepts) and links (relations) within the generated map.

Following the extraction process, an initial automatic layout is created. Leveraging techniques referenced in previous studies [de Castro et al. 2015], this layout forms a visually coherent and meaningful representation of the information contained within the text, thus producing the main artifact: an automatically generated CM.

4. Results and Implications

This section will present the preliminary results of implementing the proposed methodology, reconstructing the experiments described in [Aguiar and Cury 2017]. We draw a comparative analysis between the CMs produced by humans, the results reported by the authors, and those obtained through our approach, to highlight our methodology's effectiveness and potential within the same experimentation framework. Afterward, we critically reflect on the comparison procedure and the data obtained, both in the authors' experiments and ours, providing valuable information on the practical implications, care related to the experiment setup, and the treatment of the data obtained.

4.1. Experiment Settings

[Aguiar and Cury 2017] engaged five graduate students for their experimental process. After receiving requisite training on CM generation from text sources, these students were tasked with independently creating CMs based on a pre-selected text provided by the authors. These CMs were later juxtaposed with the outcomes yielded by the automatic generation tool proposed by the authors, using precision and recall metrics, scrutinizing individual participant results against those rendered by the automatic generator.

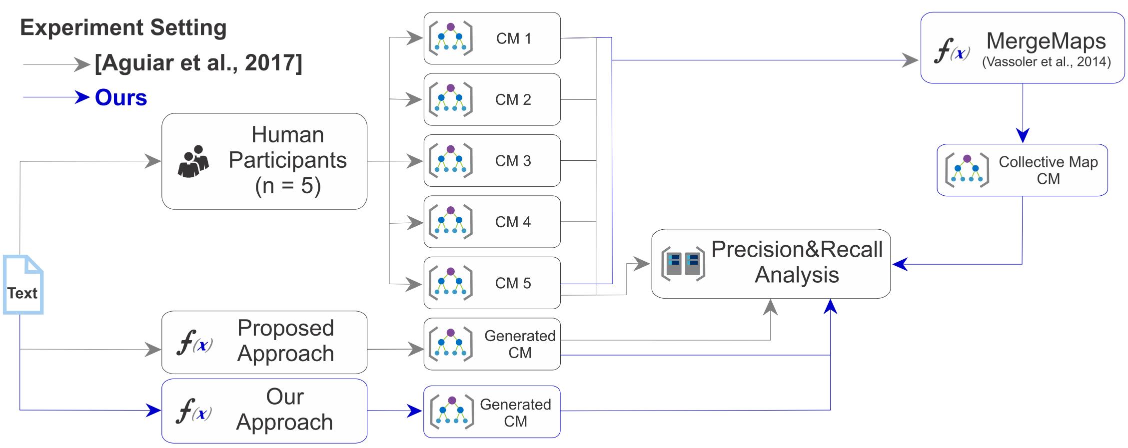

Figura 4. Comparison of Experiment Settings

In our experimental design, we implemented a slight modification. Given that our primary objective was to compare the outcomes of automatic generators against collective human effort rather than individual human performance, we amalgamated the CMs produced by the five participants. Leveraging the merging algorithm proposed by [Vassoler et al. 2014], we generated a collective CM encompassing all propositions identified by the five participants. This merged map then served as the basis for comparison against the results presented in [Aguiar and Cury 2017] and those yielded by our approach. Fig. 4 provides a detailed depiction of our experimental setup, elucidating the comparison parameters and the metrics employed.

4.2. Results

This experiment was conducted using the introductory section of the well-regarded article by [Novak and Cañas 2008] as a data source. This seminal work in Concept Mapping presented a suitable, content-rich, and complex text source for our experiment. We used the ideas and propositions presented in the introductory section to test the effectiveness of our approach in recognizing and extracting key concepts and generating a corresponding concept map. This allowed us to test our approach against the other mentioned CMs rigorously.

It starts by detailing the size of the CMs obtained in the procedures outlined in the previous subsection. The main object of this study, the MC generated by our approach, is shown in Fig. 5.

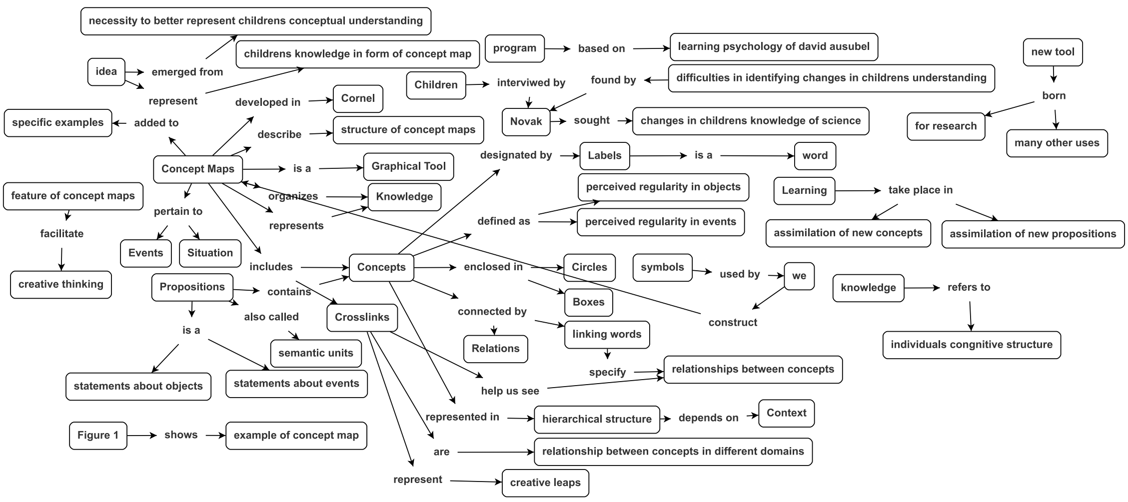

Figura 5. Generated CM from Our Approach

To encapsulate the scale of CMs produced, we provide a summary in Tab. 1, through which it is possible to observe that: the MC obtained from the fusion of the five CMs generated by human participants comprises 63 concepts and 137 propositions; CM produced by the methodology proposed by [Aguiar and Cury 2017] encapsulates 80 concepts and 123 propositions; in contrast, the CM processed by our approach contains 51 concepts and 43 propositions.

Concerning identifying concepts and propositions, as summarized in Table 2. Regarding the identification of concepts, our approach demonstrated a considerable improvement in precision by around 23.5% (0.603 versus 0.488) and a similar recall within 2.65% (0.619 versus 0.603). This result indicates that our method was more accurate

Tabela 1. Generated CMs Summary

Generated CMs Summary			
	Merged	[Aguiar et al., 2017] approach	Our approach
Number of Concepts	63	80	51
Number of Propositions	137	123	43

Tabela 2. Precision and Recall Results

		Concepts on [Aguiar et al., 2017] approach	**Concepts on Our approach**
Concepts on Merged	*precision*	0,488	**0,603**
	recall	0,603	**0,619**
		Propositions on [Aguiar et al., 2017] approach	**Propositions on Our approach**
Propositions on Merged	*precision*	0,358	**0,791**
	recall	0,307	**0,409**

in identifying true positives and minimizing false positives, thus demonstrating a higher quality of results.

As for the identification of propositions, our approach surpassed the previous method in precision by a substantial margin of approximately 121% (0.791 versus 0.358) and displayed superior recall by around 33.2% (0.409 versus 0.307). This result suggests that our approach is better in identifying true propositions from the given texts and minimizing the omission of valid propositions, thus ensuring more comprehensive coverage of the information contained in the input text.

These superior precision and recall results reflect the efficacy of our approach. The improved precision denotes a higher rate of relevant concepts and propositions identified and a lower rate of false positives. The enhanced recall illustrates that our method successfully identifies a more significant proportion of true positives. Consequently, these findings suggest that our model offers a more accurate and exhaustive method for extracting concepts and propositions from the text.

4.3. Caveats and Considerations

Our methodology employed both quantitative and qualitative analytical techniques to facilitate a comprehensive understanding of the results. In this process, we acknowledged equivalent propositions as valid. For instance, the propositions includes(concept maps, concepts) and has(concept maps, concepts) were treated as equivalent in our computations. This is an important aspect to consider when interpreting our results, as it significantly evaluates accuracy and recall.

Additionally, we evaluated the equivalence of more straightforward propositions with their more complex counterparts. As an illustrative example, the proposition definedAs(concept, perceived regularity in events) was deemed equivalent to the combined propositions definedAs(concept, regularity) and perceivedAt(regularity, event), as both structures essentially communicate the same core idea. It is important to note that the results reported in this study were obtained under controlled conditions and with limited sample size. As such, they may need to fully reflect the performance of the approach in broader, real-world contexts. In future research, we aim to expand the scope of our experiments to encompass diverse scenarios and population characteristics. This will provide us with a more nuanced understanding of the effectiveness of our approach and enable us to fine-tune it further based on these additional insights.

5. Conclusions and Future Works

This paper presents a novel approach to automatically extracting propositions from unstructured texts and generating concept maps from these propositions. Leveraging the capabilities of large language models (LLMs) like GPT-4, we devised an architecture that replaces a complex stack of traditional natural language processing (NLP) algorithms with a fine-tuned LLM. This architecture simplifies the process and demonstrates remarkable promise in proposition extraction and concept map generation.

Through a series of controlled experiments, we demonstrated our approach's superior precision and recall compared to previous methods. These results highlight the potential benefits of using LLMs for NLP tasks, particularly in education and learning tools.

However, it is important to acknowledge the limitations of our research. The results were obtained under controlled conditions with a limited sample size, which may not fully reflect real-world scenarios. Additionally, despite employing a careful analysis process that considers the equivalence of simple and complex propositions, natural language's inherent variability and subtlety can pose challenges.

Looking ahead, our future work will focus on scaling up our approach and testing it in diverse scenarios and with broader population characteristics. This will provide us with a more nuanced understanding of its strengths and potential areas for improvement. Furthermore, we intend to explore ways of integrating the proposed architecture with other educational tools and systems to enhance learning experiences. The current work is only the first step on a promising journey toward harnessing the full potential of LLMs in the educational domain.

Referências

Acharya, A. and Sinha, D. (2017). An educational data mining approach to concept map construction for web based learning. *Informatica Economica*, 21(4):41–58.

Aguiar, C. and Cury, D. (2017). Mineração de mapas conceituais a partir de textos em português. In *Brazilian Symposium on Computers in Education (Simpósio Brasileiro de Informática na Educação-SBIE)*, volume 28, page 1117.

Aguiar, C. Z., Cury, D., and Zouaq, A. (2016). Automatic construction of concept maps from texts. In *Proceedings of the 7th International Conference on Concept Mapping*, volume 2, pages 20–30, Tallinn, Estonia. CMC 2016.

Bai, S.-M. and Chen, S.-M. (2008). Automatically constructing concept maps based on fuzzy rules for adapting learning systems. *Expert Systems with Applications*, 35(1):41–49.

Baker, R. S. (2014). Educational data mining: An advance for intelligent systems in education. *IEEE Intelligent Systems*, 29(3):78–82.

Booth, D. E. (2007). Data mining methods and models. *Technometrics*, 49(4):500–500.

Chen, S.-M. and Sue, P.-J. (2013). Constructing concept maps for adaptive learning systems based on data mining techniques. *Expert Systems with Applications*, 40(7):2746–2755.

Cowan, A. M. (2002). Data mining in finance: Advances in relational and hybrid methods: Boris kovalerchuk and evgenii vityaev (eds.), kluwer academic publishers, norwell, massachusetts, 2000, hb us $120, isbn 0-7923-7804-0. *International Journal of Forecasting*, 18(1):155–156.

de Castro, R. N., Perin, W. A., and Cury, D. (2015). Layouts automáticos para mapas conceituais-um serviço integrado a uma plataforma de serviços web. In *XX Congresso Internacional de Informática Educativa (TISE)*.

de la Villa, M., Aparicio, F., Maña, M. J., and de Buenaga, M. (2012). A learning support tool with clinical cases based on concept maps and medical entity recognition. In *Proceedings of the 2012 ACM International Conference on Intelligent User Interfaces*, IUI '12, page 61–70, New York, NY, USA. Association for Computing Machinery.

Ecoffet, A. (2023). Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Hirschman, L., Park, J. C., Tsujii, J., Wong, L., and Wu, C. H. (2002). Accomplishments and challenges in literature data mining for biology. *Bioinformatics*, 18(12):1553–1561.

Lai, C.-F., Chen, S.-Y., Tsai, C.-W., Chang, Y.-C., and Su, Y.-S. (2017). Using information retrieval to construct an intelligent e-book with keyword concept map. *Eurasia Journal of Mathematics, Science and Technology Education*, 13(10):6737–6747.

Li, L., Tang, H., Wu, Z., Gong, J., Gruidl, M., Zou, J., Tockman, M., and Clark, R. A. (2004). Data mining techniques for cancer detection using serum proteomic profiling. *Artificial Intelligence in Medicine*, 32(2):71–83.

Novak, J. D. and Cañas, A. J. (2008). The theory underlying concept maps and how to construct and use them. *IHMC*.

Nugumanova, A., Mansurova, M., Alimzhanov, E., Zyryanov, D., and Apayev, K. (2015). Automatic generation of concept maps based on collection of teaching materials. In *International conference on data management technologies and applications*, volume 2, pages 248–254. SCITEPRESS.

Oppl, S. and Stary, C. (2011). Effects of a tabletop interface on the co-construction of concept maps. In *Human-Computer Interaction–INTERACT 2011: 13th IFIP TC 13 International Conference, Lisbon, Portugal, September 5-9, 2011, Proceedings, Part III 13*, pages 443–460. Springer.

Perin, W. and Cury, D. (2016). Uma plataforma de serviços para mapas conceituais. In *Brazilian Symposium on Computers in Education (Simpósio Brasileiro de Informática na Educação-SBIE)*, volume 27, page 230.

Perin, W., Cury, D., and Menezes, C. d. (2015). imap & cmpaas-de ferramenta à plataforma de serviços para mapas conceituais. *Revista Brasileira de Informática na Educação*, 24(3).

Perin, W. d. A. (2014). imap: um mecanismo de inferência para mapas conceituais. *Repositorio de Dissertacoes UFES (repositorio.ufes.br:10/1484)*.

Perin, W. d. A., Cury, D., and Menezes, C. (2014). Nlp-imap: Integrated solution based on question-answer model in natural language for an inference mechanism in concepts maps. In *Proceedings of the 14th International Conference on Concept Mapping*.

Qasim, I., Jeong, J.-W., Heu, J.-U., and Lee, D.-H. (2013). Concept map construction from text documents using affinity propagation. *Journal of Information Science*, 39(6):719–736.

Romero, C. and Ventura, S. (2007). Educational data mining: A survey from 1995 to 2005. *Expert Systems with Applications*, 33(1):135–146.

Shao, Z., Li, Y., Wang, X., Zhao, X., and Guo, Y. (2020). Research on a new automatic generation algorithm of concept map based on text analysis and association rules mining. *Journal of ambient intelligence and humanized computing*, 11:539–551.

Tseng, S.-S., Sue, P.-C., Su, J.-M., Weng, J.-F., and Tsai, W.-N. (2007). A new approach for constructing the concept map. *Computers & Education*, 49(3):691–707.

Vassoler, G. A., Perin, W. d. A., and Cury, D. (2014). Mergemaps–a computational tool for merging of concept maps. In *Proceedings of the 14th International Conference on Concept Mapping*.

Wang, M. and Hu, Y. (2016). I-raid: A novel redundant storage architecture for improving reliability, performance, and life-span of solid-state disk systems. In *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, SAC '16, page 1824–1831, New York, NY, USA. Association for Computing Machinery.

Wang, W., Cheung, C., Lee, W., and Kwok, S. (2008). Mining knowledge from natural language texts using fuzzy associated concept mapping. *Information Processing & Management*, 44(5):1707–1719.

Yang, J., Jin, H., Tang, R., Han, X., Feng, Q., Jiang, H., Yin, B., and Hu, X. (2023). Harnessing the power of llms in practice: A survey on chatgpt and beyond. *arXiv preprint arXiv:2304.13712*.

Zouaq, A. and Nkambou, R. (2009). Evaluating the generation of domain ontologies in the knowledge puzzle project. *IEEE Transactions on Knowledge and Data Engineering*, 21(11):1559–1572.

Žubrinić, K., Obradović, I., and Sjekavica, T. (2015). Implementation of method for generating concept map from unstructured text in the croatian language. In *2015 23rd International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, pages 220–223.