A Topic–Word Graph for Topic Modeling Visualization

Ana Rita Peixoto

ISTAR, Instituto Universitário de Lisboa (ISCTE-IUL), Portugal

Abstract

Topic Modeling (TM) is a widely used method for discovering hidden themes in unlabeled text collections, but the resulting topics are often difficult to interpret. Most existing visualizations present topics separately and do not show the words that link them, offering only a partial view of how topics relate to one another. This work introduces a graph-based representation in which topics and words are modeled as connected nodes. This structure makes shared words explicit and allows users to observe how individual terms contribute to multiple thematic areas within the corpus. By encoding topic–word relationships as edges, the visualization reveals the distribution of these connections and provides a more detailed perspective than traditional topic-centric displays. The graph includes four node types, defined by their role and degree (number of connections). The developed tool, `topicwordgraph`, offers two complementary visualizations: a main graph representation that distinguishes node types by color, and two bar charts that summarize topic overlap and word memberships. Together, these views help users understand how topics intertwine through shared vocabulary. By making the shared words between topics explicit, this tool offers a practical and effective way to interpret how topics are interconnected.

2012 ACM Subject Classification Computing methodologies → Natural language processing

Keywords and phrases Topic Modeling, Data Visualization, Graph-Based Representation, Topic Interpretability, Text Mining

Digital Object Identifier 10.4230/OASIcs.SLATE.2026.6

Supplementary Material *Software (Source Code)*: `https://github.com/ritapx/topicwordgraph` [14]; archived at `swh:1:dir:2f09eb9f388de51f82caa9f5e7a78d537a07314a`

Funding This work was partially supported by national funds through Fundação para a Ciência e a Tecnologia, I.P. (FCT), ISTAR-Iscte Pluriannual Projects UID/04466/2025 (`https://doi.org/10.54499/UID/04466/2025`).

1 Introduction

When there is a need to discover themes within text corpora, independent of the research domain, the primary technique used is Topic Modeling (TM) [3]. In this framework, the model outputs a list of the top-k words associated with each topic, along with the corresponding topic-word distribution. This list is used both to describe the topics and to evaluate the model's overall performance and interpretability. However, top-k word lists may not always provide sufficient semantic clarity, as they can hide topic coherence and delay human interpretability [9]. As a result, newer studies have emphasized the need for complementary evaluation strategies that better capture the nuanced structure of topic–word relationships [1], highlighting how central these connections are for understanding what each topic truly represents and how they can be connected.

Despite their computational simplicity, topic models often present interpretational and communicational challenges when reduced to word lists and word-cloud visualizations [5, 11]. Some visualizations have been developed to support understanding of topic models and to assess whether topics are well separated, such as the well-known LDAvis tool [17] applied to the Latent Dirichlet Allocation (LDA) algorithm [4]. Nevertheless, these approaches tend to be more effective for experts in the field than for communicating results to a broader audience [3]. Highlighting the need to develop complementary visualization strategies that bridge that gap and capture the topic-word connections.

The visualization of topics and their connecting words remains largely unexplored. Most existing approaches focus on isolated topic representations rather than on shared vocabulary. Although an ideal topic model would exhibit no overlap in their top words, this assumption rarely holds in practice, where topic boundaries tend to blur, and word overlap naturally emerges [1]. In practice, word overlap across topics reflects the inherent interconnections among them and often reveals a central theme that summarizes the source corpus. With this work, we aim to address this gap by answering the following research question:

- How can topic models be visualized to effectively reveal and interpret the interconnection of topics through shared words?

To address this research question, a graph-based visualization was developed and implemented in Python - `topicwordgraph`[1]. The objective of this tool is to provide a visual aid for how topics intertwine, to uncover the specific words that connect them, and to provide a clearer representation of the thematic and semantic structure underlying the data.

2 Related Work

Identifying meaningful patterns and extracting themes is a common need across various domains, and TM can perform it effectively [3]. Recent studies use it for various tasks, from a complementary tool in systematic literature reviews [7] to its application in applied genomic science [20].

Topic Modeling Visualization

Most studies, when interpreting topics, use word clouds and tables or bar charts with the top words for each topic. The following is a brief description of some TM visualization tools that aim to improve that examination.

LDAvis is a combined visualization that aims to increase the interpretability of each topic, understand how prevalent each topic is, and how topics relate to each other [17]. For that, it offers two visualizations: a global view of the topic model and a bar chart showing the words most useful for interpreting the selected topic. The global view of the topic model is a 2D plane, with topics represented as circles; the distance between circles indicates similarity, and their sizes indicate prevalence.

A recent study proposes an integrated application called WordMap [18]. The authors explain that most tools are segmented, offering a single tool with various text mining techniques. In terms of visualization, this tool provides: word clouds, word distributions over topics in a bar chart, an inter-topic distance map (equivalent to LDAvis), topic distributions over the corpus, a hierarchical clustering graphic with semantic associations, and a heat map for visualizing semantic similarity.

Similar to the WordMap, the topicwizard was designed to enhance topic visualization [11]. It also incorporates an inter-topic distance map, topic-word distributions, and word clouds. However, it differs by presenting a word map that enables interactive analysis of the corpus's semantic landscape using word embeddings, and by providing a document view exploratory analysis.

[1] `https://github.com/ritapx/topicwordgraph`

However, these tools are mostly used for the interpretation and evaluation of topics by field experts. To the best of our knowledge, there is a gap in the communication of the extracted topics. As well as a gap in the analysis of the shared words by topic as a semantic connection between the topics, and as a description of the data domain used in the TM.

3 Topic-Word Graph

A Python tool was created for TM visualization, the `topicwordgraph`, enabling the extraction of information about the words that define each topic and their interconnections. It transforms a topic model into a network with the structure of the Topic-Word graph.

This section begins by formally specifying the Topic-Word graph using the TM components, the topics (T), and the words (W). Then it explains how to interpret a network obtained by applying a specific model to the Topic-Word graph. Thirdly, it describes how visualization elements are utilized to facilitate the interpretation of the resulting graph. Finally, it briefly explains the implementation and how to use the developed tool.

3.1 The Graph

TM produces a collection of n topics, which we denote by $T = \{T_1, T_2, ..., T_n\}$. Each topic T_i is represented by its top-k words, $T_i = \{w_1, w_2, ..., w_k\}$, these are the words most strongly associated with the topic according to the model. By taking the union of all topic words, we obtain the set W, which contains every distinct word that appears in at least one topic: $W = \bigcup_{i=1}^{n} T_i$. This set represents the vocabulary used to describe the topics.

A graph is a mathematical structure defined as a pair $G = (V, E)$, where V is a set of nodes and $E \subseteq V \times V$ is a set of edges that represent relationships between nodes [6]. To define a topic model as a graph, we represent the nodes as topics or words, and the model's relationships as edges, specifically, the membership of a word in a topic.

The node set is $V = T \cup W$, meaning that every topic $T_i \in T$ and every word $w \in W$ is a node. These two types of nodes form two disjoint subsets. The set of edges, E, is the set of connections between words and topics. Each edge means that a word belongs to a topic, resulting in a bipartite graph where $E \subseteq T \times W$. No edges occur between two topics or between two words. This restriction creates a bipartite graph, with topics on one side and words on the other. The edges represent the relation "word belongs to topic," which does not require directionality for analysis. Additionally, the edges do not need a weight assigned since they represent a single membership.

The Topic-Word graph, G, is therefore an undirected, unweighted, and bipartite graph that can be represented as follows: $G = (T \cup W, E), E = \{(T_i, w) | w \in T_i\}$. Providing a structural representation of the topic model, where topics and words are nodes, and edges encode the membership relation between them. Figure 1 illustrates the Topic-Word graph.

3.2 The Interpretation

Once defined, the Topic–Word Graph enables the use of graph-theoretic characteristics to analyze its structure and meaning. These will provide valuable insight into how topics relate to one another, how words are distributed across topics, and how the vocabulary contributes to the overall coherence of the topic model. Allowing the extraction of new information that would otherwise be harder to conclude.

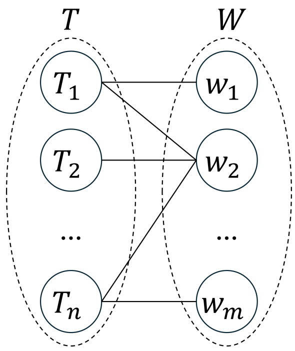

Figure 1 Topic-Word Graph illustration, where T is the set of the n topics in the topic model, W the union of the top-k words of each topic (resulting in m different words) and $G = (T \cup W, E), E = \{(T_i, w) | w \in T_i\}$, a undirected, unweighted and bipartite graph ($E \subseteq T \times W$).

The Topic–Word Graph naturally facilitates the analysis of words that appear across multiple topics. The graph measure employed to know which word belongs to one or more topics is the node degree. The degree of a node is the sum of the edges linked to it [2]. Since it is a bipartite graph, there are no word nodes linked to other word nodes, so the degree value indicates the number of topics the word belongs to.

The degree of a word node can vary from one to n, the number of topics. When the degree of a word is one, it means that the word only belongs to a topic, and when it is n, it belongs to all the topics. Additionally, the degree of all topic nodes is k, the number of top words used to create the graph. For instance, in Figure 1, the degree of the word node w_1 is one, indicating that it belongs to only one topic. And the degree of the word w_2 is three, belonging to three topics. We define a word node with a degree of one, a *singleton word*, because it belongs to only one topic. In contrast, a *central word* is a word node with a degree value greater than one, meaning it belongs to more than one topic. When a *central word* has degree n, it belongs to all topics and is a special case; it is called a *global central word*.

The *singleton words* (degree $= 1$) appear exclusively in a single topic and therefore contribute uniquely to the semantic interpretation of that topic. Since *singleton words* do not connect across topics, they play no role in inter-topic relations and instead work as a reinforcement of the distinctiveness of individual topics. In a coherent topic model, each topic should be unique and specific [13].

In contrast, *central words* (degree > 1) act as connectors between topics. Understanding *central words* is important because they reveal the shared theme across topics. When those words appear in all the topics, with a degree value of n, the *global central words* describe and summarize the data. Furthermore, if a model has mostly *central words*, they can reveal semantic overlap between the topics, suggesting that those topics may not be the most suitable [13, 16].

3.3 The Visualization

The graph visualization incorporates visual elements to enable the interpretation explained. In terms of format, circles represent the graph's nodes, and in the middle is written the name of each topic or the specific word. Topic nodes are larger than word nodes to help separate the two concepts. Lastly, to draw the nodes in a dispersed, non-overlapping layout, the `networkx` library's `spring_layout` was used. Additionally, to allow the user to change the layout as desired and improve the communication of results, the graph is interactive, enabling them to move the nodes. The visualizations were implemented using two Python libraries: `networkx` [8] and `matplotlib` [10].

Color is the key element in this visualization, enabling immediate understanding of the four node types: *topics*, *singleton words*, *central words*, and *global central words*. We defined a default palette, the one applied in the visualizations of the results section above. The default palette is: blue for the *topics* and a lighter blue for the *singleton nodes*, since they characterize the topics; and green for the *central words*, with a darker shade of green for the *global central words* to create an emphasis. We choose this approach to achieve a comprehensive analysis.

Lastly, the `topicwordgraph` tool allows the creation of two bar charts as a complementary visualization of the graph. A bar chart named "Topics with most overlapping words", which shows the count of the overlapping words for each topic name. Enabling easy recognition of the topics with higher or lower levels of interconnection. Additionally, a second bar chart named "Words with degree > 1", which displays the *central words* and their degree values. Facilitating the distinction between them, effortlessly identifying which ones are the ones that create the interconnection between the topics.

3.4 The Tool Usage

To use the `topicwordgraph` tool, the user needs to provide the topic names and the corresponding top-k words in a dictionary format where the key is the topic name, and the value is a list of words. The following is a simple example of how to use the tool to create a graph visualization (`.plot()`) and the bar charts (`.plot_bar_charts()`).

```
from topicwordgraph import TopicWordGraph

topics = {
    "Topic Name 1": ["word1", "word2", "word3"],
    "Topic Name 2": ["word1", "word4", "word5"],
    "Topic Name 3": ["word2", "word6", "word7"]
}

twg = TopicWordGraph.from_named_topics(topics)

# Interactive Topic-Word graph
twg.plot()

# Bar charts
twg.plot_bar_charts()
```

4 Results

The purpose of this Section is to showcase visualizations created with the `topicwordgraph` tool, emphasizing its advantages and limitations. For that, three different TM were used to create the visualizations: one from a prompt in a large language model, the GPT-5.3 [12]; another from the results of a literature paper that applied a TM to startups' tweets [15]; and, lastly, one from a classic example with the data of the 20 newsgroups dataset [19].

Simulation with GPT-5.3

It was important to create a visualization that illustrates all the components, including the four node types that enable the understanding of the interconnections in a topic model. Since most real-world applications may not allow us to achieve that goal, we simulate a topic model using a prompt in ChatGPT with the GPT-5.3 version. In the prompt, it was assured that all node types would appear in the result.

The following is the prompt used: *Please simulate a topic modeling and provide a list of five topics and their top 10 words. One or two words must belong to the five topics; other words must belong to more than one topic, but not all topics. You must have words that are common to more than one topic, and words that belong only to one topic. The topic model result must be a good model, and the domain should be data visualization. The format of the result must be: "topics = { "Topic name 1": ["word1", "word2", ...], "Topic name 2": ["word3", "word 4", ...] }".*

The resultant list was then applied to the graph to create the visualizations. Figure 2 shows the graph, and the corresponding bar charts are in Figure 3.

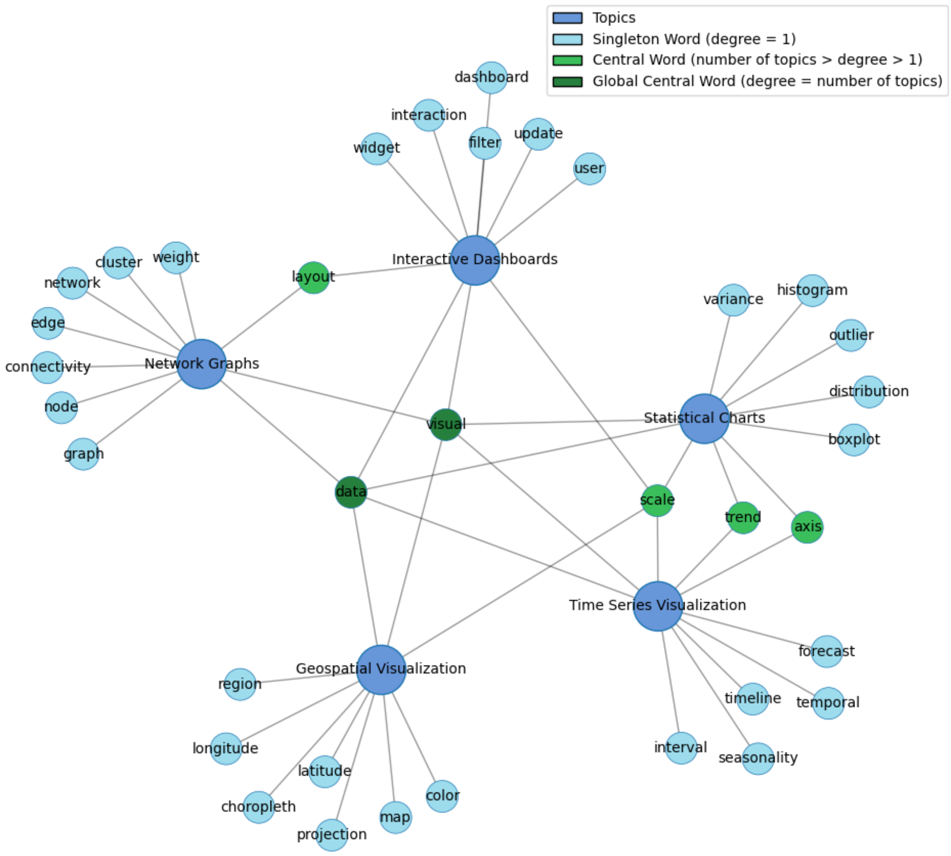

Figure 2 Topic-word graph visualization of the topic model created by ChatGPT (GPT-5.3) about data visualization.

The graph (Figure 2) shows the five topics: "Interactive Dashboards", "Network Graphs", "Geospatial Visualization", "Statistical Charts", and "Time Series Visualization". The topics "Statistical Charts" and "Time Series Visualization" have the highest word overlap. The

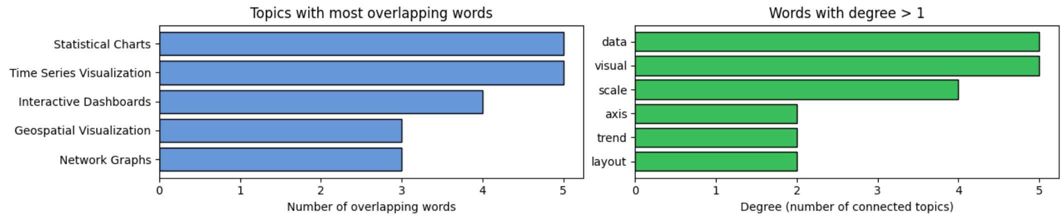

Figure 3 Bar charts of the topic model created by ChatGPT (GPT-5.3) about data visualization.

topic "Interactive Dashboards" is connected to "Network Graphs" via the word "layout" and to "Statistical Charts" via the word "scale". The visualization shows that the words "data" and "visualization" are central to the model's domain, connecting all the topics and revealing the shared theme across them. The bar charts, in Figure 3, complement those findings.

From the visualizations, it is possible to extract that the model is about data visualization, using the *global central words* "data" and "visual", further "layout", "scale", "trend", and "axis" are valuable concepts in that domain, since they connect other topics.

Topics of Startups' Tweets

The following is a scenario of TM applied to real data. The authors in [15] used TM to understand what Portuguese IT startups posted on Twitter from 2015 to 2020, using a dataset of 15,557 tweets from eight startups. The results comprise a coherent model (using LDA) with five topics described by their 10 most important words, which enabled the application of the results to a Topic-Word graph.

One of the conclusions was that startups posted about the following five topics: "Fintech and ML", "IT", "Bank and funding", "Business operations", and "Product / service RD". Additionally, the authors performed a time-based analysis, discovering that those themes predominance change along the startup's maturity. These are valuable findings, but applying graph visualization to the results may enable us to extract more knowledge. Figure 4 shows the resultant graph, and the corresponding bar charts are in Figure 5.

From the visualizations, it is possible to see that the words "lisbon", "learn", "webinar", "machinelearning", "cloud", "service", and "read" are *central words*. Together, these words characterize the domain of Portuguese IT startups, highlight recurring themes in their content, and reveal what these startups always focus on. This analysis could enrich the work already done.

Topics of the 20 Newsgroup Dataset

The 20 newsgroup dataset is widely used and comprises approximately 20,000 newsgroup documents. In [19], the authors applied various TM algorithms to the dataset and evaluated the results using annotators. The evaluation was made based on the top-10 most important words for the topics. The MRF-LDA achieved the best model with four topics: "Crime", "Sex", "Sports", and "Health". The list of top words and the topic names enabled the use of the `topicwordgraph` tool. The resulting visualization of applying the top-10 words for each topic to the Topic-Word graph is in Figure 6, and Figure 7 displays the complementary bar chart.

The graph visualization provides an easy overview of the four topics: two are completely independent, with only *singleton words*, and one *central word* connects the other two. Since only one of the 40 words belongs to more than one topic, this implies that those four topics

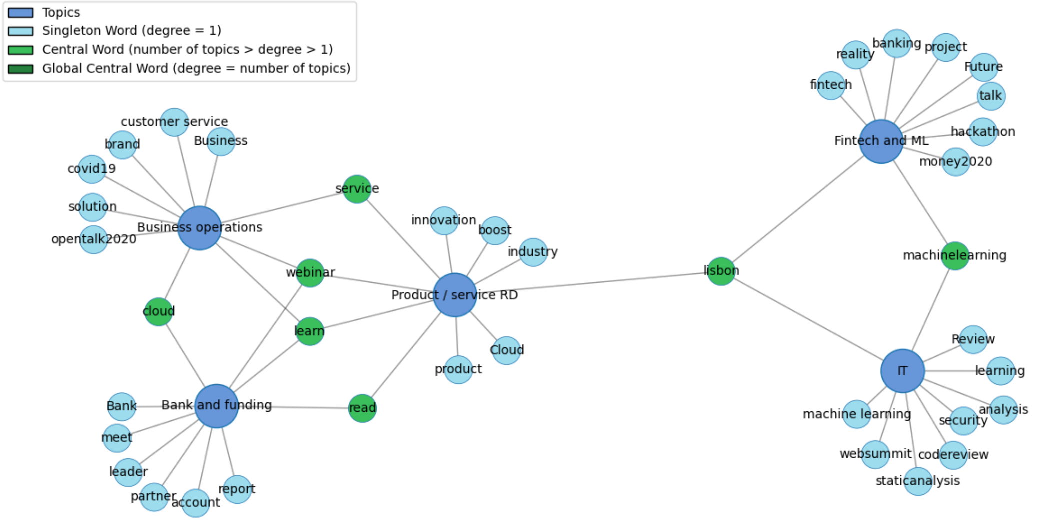

Figure 4 Topic-word graph visualization of the topic model about what Portuguese IT startups posted on Twitter [15].

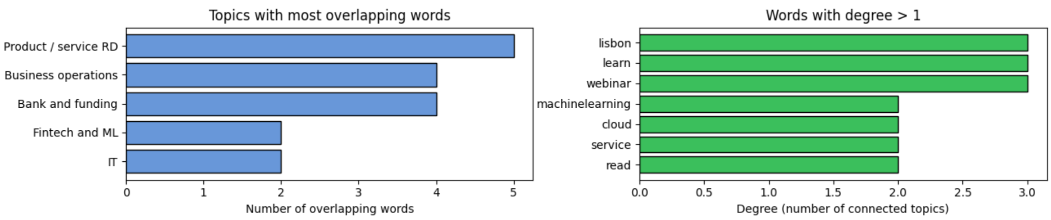

Figure 5 Bar charts of the topic model about what Portuguese IT startups posted on Twitter [15].

are very well separated. The *central word* is "child", connecting the topics "Crime" and "Sex", suggesting that some news combine those two topics with children involved. This analysis, enabled by the visualization, reveals how the newsgroup topics can be interconnected and allows a clear visualization of theme separations across the other topics.

5 Conclusions

The TM representation as a Topic-Word graph, using the developed tool, enables the visualization of: 1) how separated the topics are; 2) how some *central words* thematically connect some topics; and 3) how *central words* characterize and summarize the data domain. These three key points outline a comprehensive analysis of the TM resulting models. By representing topics and their shared words as a connected network, based on the Topic-Word graph, it is possible to effectively reveal and interpret the interconnections among topics, directly answering the proposed research question.

The main purpose of this tool's visualizations is to communicate TM results, complementing more typical visualizations, such as word clouds, with topic-word connections. However, it can also be applied to evaluate whether topics are sufficiently separated and, at the same time, whether the extracted topics well represent the data. Allowing it to be an intermediate visualization while fine-tuning a model or to provide to annotators.

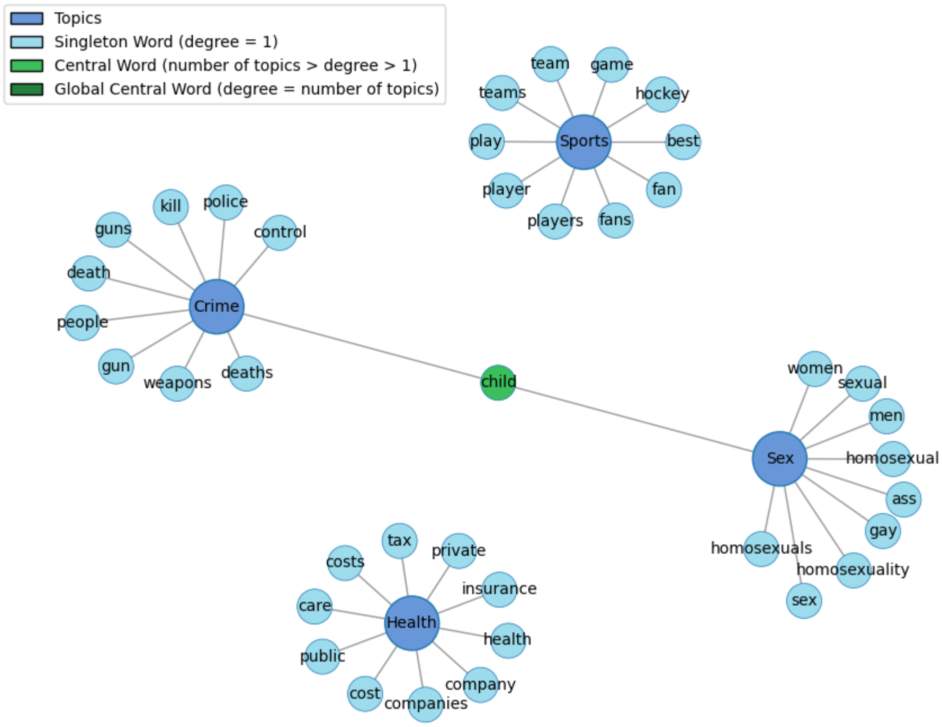

Figure 6 Topic-word graph visualization of the topic model applied to the 20 newsgroup dataset [19].

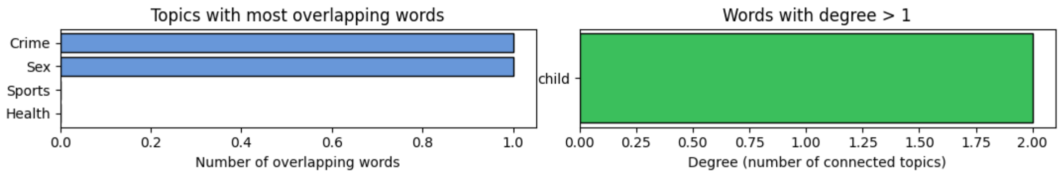

Figure 7 Bar charts of the topic model applied to the 20 newsgroup dataset [19].

Limitations

The proposed graph-based visualization has some limitations. One of them is scalability and size, since topic models with more topics produce more words, increasing visual complexity. Secondly, this work assumes that shared words imply semantic connections between topics. However, some words may appear across topics due to their high frequency, even after suitable preprocessing and the removal of stopwords. Lastly, it does not allow hierarchical connections; the graph shows the relation between topic and word, and does not include relationships with subtopics. Restricting the visualization to non-hierarchical topic models.

Future Work

Several directions can extend and strengthen this work. First, future research could focus on improving the scalability of the visualization by incorporating some filtering, such as edge-reduction techniques, to prevent visual clutter and maintain interpretability. Second, the exploration of more sophisticated methods for identifying meaningful shared words, as the current work focuses only on degree values. Third, enable hierarchy among topics, expanding the Topic-Word graph to include a new type of node: subtopics. Lastly, another promising direction is to incorporate TM evaluation metrics directly into the visualization, enabling effortless assessment of model quality.

References

1 Ana Amaro and Fernando Bacao. Topic modeling: A consistent framework for comparative studies. *Emerging Science Journal*, 8:125–139, February 2024. doi:10.28991/ESJ-2024-08-01-09.

2 Albert-László Barabási. *Network Science*. Cambridge University Press, 2016. URL: http://networksciencebook.com/.

3 Monika Bednarek. Topic modelling in corpus-based discourse analysis: Uses and critiques. *Discourse Studies*, 27:659–671, August 2025. doi:10.1177/14614456241293075.

4 David M Blei, Andrew Y Ng, and Jordan@cs Berkeley Edu. Latent dirichlet allocation michael i. jordan. *Journal of Machine Learning Research*, 3:993–1022, 2003.

5 Jonathan Chang, Jordan Boyd-Graber, Sean Gerrish, Chong Wang, and David M. Blei. Reading tea leaves: how humans interpret topic models. In *Proceedings of the 23rd International Conference on Neural Information Processing Systems*, NIPS'09, pages 288–296, Red Hook, NY, USA, 2009. Curran Associates Inc. URL: https://proceedings.neurips.cc/paper/2009/hash/f92586a25bb3145facd64ab20fd554ff-Abstract.html.

6 Reinhard Diestel. *Graph Theory*. Springer, 5 edition, 2017. doi:10.1007/978-3-662-53622-3.

7 Priyanka Gupta, Bosheng Ding, Chong Guan, and Ding Ding. Generative ai: A systematic review using topic modelling techniques. *Data and Information Management*, 8, June 2024. doi:10.1016/j.dim.2024.100066.

8 Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. Exploring network structure, dynamics, and function using networkx. *Proceedings of the 7th Python in Science Conference (SciPy2008)*, pages 11–15, 2008.

9 Alexander Hoyle, Pranav Goel, Denis Peskov, Andrew Hian-Cheong, Jordan Boyd-Graber, and Philip Resnik. Is automated topic model evaluation broken? the incoherence of coherence. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NIPS '21, Red Hook, NY, USA, 2021. Curran Associates Inc.

10 J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. doi:10.1109/MCSE.2007.55.

11 Márton Kardos, Kenneth C. Enevoldsen, and Kristoffer Laigaard Nielbo. topicwizard – a modern, model-agnostic framework for topic model visualization and interpretation. In *Proceedings of the 8th International Conference on Natural Language and Speech Processing (ICNLSP-2025)*, May 2025. doi:10.48550/arXiv.2505.13034.

12 OpenAI. GPT-5.3 (ChatGPT). https://chat.openai.com/, 2026. Large language model.

13 Derek O'Callaghan, Derek Greene, Joe Carthy, and Pádraig Cunningham. An analysis of the coherence of descriptors in topic modeling. *Expert Systems with Applications*, 42(13):5645–5657, 2015. doi:10.1016/j.eswa.2015.02.055.

14 Ana Rita Peixoto. topicwordgraph. Software, version 1.0., swhId: swh:1:dir:2f09eb9f388de51f82caa9f5e7a78d537a07314a (visited on 2026-07-17). URL: https://github.com/ritapx/topicwordgraph, doi:10.4230/artifacts.27313.

15 Ana Rita Peixoto, Ana de Almeida, Nuno António, Fernando Batista, and Ricardo Ribeiro. Diachronic profile of startup companies through social media. *Social Network Analysis and Mining*, 13:1–18, 2023. doi:10.1007/s13278-023-01055-2.

16 Emil Rijcken, Kalliopi Zervanou, Pablo Mosteiro, Floortje Scheepers, Marco Spruit, and Uzay Kaymak. Topic specificity: A descriptive metric for algorithm selection and finding the right number of topics. *Natural Language Processing Journal*, 8:100082, 2024. doi:10.1016/j.nlp.2024.100082.

17 Carson Sievert and Kenneth E Shirley. Ldavis: A method for visualizing and interpreting topics. In *Proceedings of the Workshop on Interactive Language Learning, Visualization, and Interfaces*, pages 63–70, 2014.

18 Zhijian Wei, Huiwen Zou, Patrick Cheong Iao Pang, Penny Wong On Chao, and Benjamin K. Ng. Wordmap: Text mining application of enhanced corpus segmentation and semantic topic recognition. *Applied Sciences (Switzerland)*, 15, June 2025. doi:10.3390/app15126632.

19 Pengtao Xie, Diyi Yang, and Eric Xing. Incorporating word correlation knowledge into topic modeling. In Rada Mihalcea, Joyce Chai, and Anoop Sarkar, editors, *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 725–734, Denver, Colorado, May–June 2015. Association for Computational Linguistics. doi:10.3115/v1/N15-1074.

20 Yidong Zhang, Xilin Jiang, Alexander J. Mentzer, Gil McVean, and Gerton Lunter. Topic modeling identifies novel genetic loci associated with multimorbidities in uk biobank. *Cell Genomics*, 3, August 2023. doi:10.1016/j.xgen.2023.100371.