

# SocialVisTUM: An Interactive Visualization Toolkit for Correlated Neural Topic Models on Social Media Opinion Mining

Gerhard Hagerer<sup>2</sup>, Martin Kirchhoff<sup>3</sup>, Hannah Danner<sup>2</sup>, Robert Pesch<sup>3</sup>  
Mainak Ghosh<sup>1</sup>, Archishman Roy<sup>2</sup>, Jiayi Zhao<sup>2</sup>, Georg Groh<sup>2</sup>

<sup>1</sup>Max Planck Institute for Innovation and Competition

<sup>2</sup>Technical University of Munich

<sup>3</sup>inovex GmbH

{gerhard.hagerer, hannah.danner, archishman.roy, jiayi.zhao}@tum.de, grohg@mytum.de

{martin.kirchhoff, robert.pesch}@inovex.de, mainak.ghosh@ip.mpg.de

## Abstract

Recent research in opinion mining proposed word embedding-based topic modeling methods that provide superior coherence compared to traditional topic modeling. In this paper, we demonstrate how these methods can be used to display correlated topic models on social media texts using SocialVisTUM, our proposed interactive visualization toolkit. It displays a graph with topics as nodes and their correlations as edges. Further details are displayed interactively to support the exploration of large text collections, e.g., representative words and sentences of topics, topic and sentiment distributions, hierarchical topic clustering, and customizable, predefined topic labels. The toolkit optimizes automatically on custom data for optimal coherence. We show a working instance of the toolkit on data crawled from English social media discussions about organic food consumption. The visualization confirms findings of a qualitative consumer research study. SocialVisTUM and its training procedures are accessible online<sup>1</sup>.




Figure 1: SocialVisTUM applied to our use case *organic food* - The topics, their occurrence (in brackets) and respective correlations.

## 1 Introduction

Web sources, such as social networks, internet forums, and customer reviews from online shops, provide large amounts of unstructured text data. Along with the steady development of new platforms and the increasing number of internet users, the interest in methods that automatically extract the expressed opinions along with the corresponding topics and sentiments in text data has increased in recent years. Scholars and organizations from different fields can utilize such methods to identify patterns and generate new insights. Examples are opinion researchers investigating current opinions on political and societal issues, consumer researchers interested in

consumers' beliefs about the consumption and production of goods (Danner et al., 2020), and marketing managers curious about the public perception of their products and services (Berger et al., 2020; Murphy et al., 2014). (Kirchhoff, 2019)

These domain-specific use cases are of interest for research disciplines which taken by itself are not directly related to natural language processing (NLP). Consequentially, there is a constant need to provide state-of-the-art NLP methods such that domain researchers from other fields can take advantage of them. The requirements therefore are simple usage, automatic hyperparameter optimization, minimal effort for manual labeling of text data, and built-in visualizations to give an abstract overview of the discussed topics and their relation

<sup>1</sup><https://github.com/ghagerer/SocialVisTum>

with each other.

While these practical requirements are important for domain experts, modern opinion mining approaches target specific machine learning objectives. Recently, there is a trend towards unsupervised neural methods for opinion target detection. Attention-based aspect extraction (ABAE) enables clustering of short review texts with significantly higher coherence as traditional LDA-based topic modeling, and it gives 70% F1 score for classification (He et al., 2017). This is improved recently (Karamanolakis et al., 2019; Angelidis and Lapata, 2018; Luo et al., 2019), which underlines the recent impact and potential of related techniques.

However, these have not been utilized for visualizations based on correlated topic modeling (Blei and Lafferty, 2006), where all pairs of topics "are" analyzed to determine if two topics generally tend to occur in the same texts of a given dataset. Thus, the similarity between topics can be defined. This is successfully used to connect topics (nodes) among each other based on their correlations (edges) leading to more abstract and more meaningful meta topics (graph-clusters) which additionally improves topic coherence. Consequentially, these meta topics, e.g., company-related events or research sub-disciplines (Liu et al., 2014; Maiya and Rolfe, 2014), can be successfully identified by graph-based visualization techniques. However, there is a lack of related prototypes on texts discussing consumption related issues in product reviews or social media. To the best of our knowledge, there is also no related integration of sentiment analysis into a system available for potential end users, i.e., domain experts. As according text data from customers is available on a large scale in social media, this can be considered as a shortcoming in the field.

To address all denoted issues, we propose the *SocialVisTUM* toolkit, a new visualization and labeling tool to give users a comprehensible overview of the topics discussed in social media texts. It integrates a neural method for unsupervised sentence and comment clustering based on word vectors and attention. We denote the respective clusters as topics hereafter. In addition, we provide a graph-based visualization showing the topics as labeled nodes and the correlation between them as edges. A force-directed graph layout maintains readability even while many relevant topics and topic relations are displayed. (Kirchhoff, 2019)

In our interactive graphical user interface, the

number of topics displayed and the correlation threshold required to display a connection between two topics can be dynamically adjusted. Further, contextual topic information is provided, such as the number of respective topic occurrences in the social media texts as node diameter, the correlation between the topic occurrences as edge width, example sentences from the data for each topic, a list of representative words for each topic, and the regarding sentiment distribution of a topic. It is a common practice to represent topics merely by word lists (Blei et al., 2003; Chen et al., 2014), which tend to be insufficient to comprehensively express a topic on our given dataset. (Kirchhoff, 2019)

To avoid manual labeling and to give users an immediate impression of each topic, topic labels are generated automatically based on a custom algorithm utilizing the most common WordNet hypernym in a topic's top words. Furthermore, we find that topic hypernym statistics can serve as a metric for automatic hyperparameter optimization, which in our case gives practical advantages over widely used coherence scoring metrics.

In addition to a more detailed description of our *SocialVisTUM* toolkit, we show the results of a case study based on social media texts from online commenters debating about organic food consumption. We demonstrate that the correlated topics give a meaningful graph representation of the social media discussions supporting the understanding of the concerns of consumers. In this regard, we also show how the combined illustration of different types of relevant topic and sentiment information and automatic labeling of clusters are a contribution.

## 2 Related Work

Correlated topic models were introduced 2006 (Blei and Lafferty, 2006; Li and McCallum, 2006) to improve topic coherence and to provide graph visualizations based on topics as nodes and their correlations as edges. This shows potential to improve text mining for the end user as "powerful means of exploring, characterizing, and summarizing large collections of unstructured text documents" (Maiya and Rolfe, 2014). Meta topics, such as research domains and their inter-disciplinary overlaps, can thus be described clearly, automatically, and empirically (Blei and Lafferty, 2007).

These correlated topic models are applied for

more sophisticated visualization approaches. *TopicPanorama* models technology-related topics from various text corpora, including newspaper articles, blogs, and tweets (Liu et al., 2014). Here, the domain expert is given the option to interactively modify the matching result of the labeled topic graph. Another topic visualization called *topic similarity networks* is particularly addressing the visualization of large document sets (Maiya and Rolfe, 2014). While claiming good scalability regarding the number of documents, beneficial methods to achieve automatic topic labeling are successfully quantified. *TopicAtlas* provides a graphical user interface to explore text networks, such as hyperlinked webpages and academic citation networks. For manual mining purposes, topic models are generated and related to one another to facilitate manual navigation and finding of relevant documents (He et al., 2016). These examples show a steady, meaningful, and promising development regarding the visualization of correlated topic modeling, partially also applied to social media texts such as micro-blogs. However, these examples do not include sentiment analysis as means to conduct market research and quantify customer satisfaction in specific and not yet explored market domains. Furthermore, the widely used latent Dirichlet allocation (LDA) technique tends to be incoherent on short texts, such as, product reviews or social media comments, and thus insufficient to detect opinion targets in an unsupervised manner (He et al., 2017).

Automatic topic coherence optimization can be seen as desirable for a topic modeling visualization toolkit such as SocialVisTUM, which tries to minimize manual optimization efforts for non-technical users. Therefore, we refer to two widely used coherence definitions (Ghosh, 2020). Firstly, word co-occurrence-based methods measure how often pairs of representative topic words co-occur in the training data set or in an external reference data set. In that regard, it has been shown that the evaluation methods UMass, UCI and NPMI correlate with human judges (Stevens et al., 2012; Newman et al., 2010; Mimno et al., 2011; Bouma, 2009; Ding et al., 2018) and are considered to be a default metric for topic coherence. Secondly, word embedding similarity based coherence scores are recently utilized as these are also based on word co-occurrence statistics (Pennington et al., 2014) and behave similar to NPMI coherence scoring (Ding et al., 2018), resulting in high correlation with hu-

man perception. These methodologies show the undesirable effect of no distinct local optimum when the hyperparameters of the models are changed, e.g., number of clusters or vocabulary size. On our data, these parameters increase together with the coherence scores, while the subjective performance, i.e., the human perception, actually decreases. We describe this effect and our solution in the case study section.

### 3 Clustering Architecture

The unsupervised neural network model called attention-based aspect extraction (ABAE) (He et al., 2017) clusters sentences based on GloVe word embeddings (Pennington et al., 2014) and attention (Bahdanau et al., 2014) to focus on the most important words in a sentence. Every sentence  $s$  is represented by a vector  $z_s$  that is defined as the weighted average of all the word vectors of that sentence. The weights are attentions calculated based on the contribution of the respective words to the meaning of the sentence and the relevance to the topics. These topics are defined by the actual centroids. In their publication, the topics are mapped to aspect classes for unsupervised aspect extraction, which we do not do for our case. (Kirchhoff, 2019)

The topics are initialized as the resulting centroids of k-means clustering on the word embeddings of the corpus dictionary. These are then stacked as topic embedding matrix  $\mathbf{T}$ . During training, ABAE calculates sentence reconstructions  $\mathbf{r}_s$  for each sentence. These are linear combinations of the topic embeddings from  $\mathbf{T}$  and defined as

$$\mathbf{r}_s = \mathbf{T}^\top \cdot \mathbf{p}_t, \quad (1)$$

where  $\mathbf{p}_t$  is the weight vector over  $K$  topic embeddings. Each weight corresponds to the probability that the input sentence belongs to the associated topic.  $\mathbf{p}_t$  is obtained by reducing the dimension of  $\mathbf{z}_s$  to the number of topics  $K$  and applying softmax such that

$$\mathbf{p}_t = \text{softmax}(\mathbf{W} \cdot \mathbf{z}_s + \mathbf{b}), \quad (2)$$

where  $\mathbf{W}$  and  $\mathbf{b}$  are trainable and matrix weights and a bias vector respectively. The topic embeddings  $T$  are updated during training to minimize the reconstruction error  $J(\theta)$  between  $\mathbf{r}_s$  and  $\mathbf{z}_s$  based on the contrastive max-margin objective function. Since words and topics share the same dimensionality, cosine similarity between both can be used to

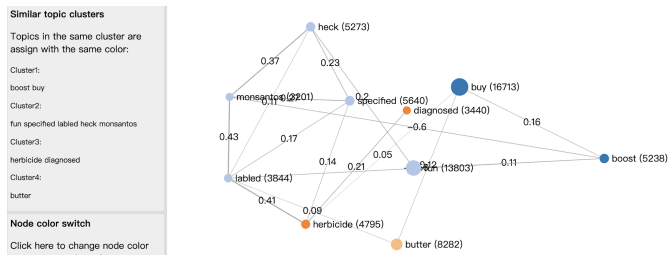


Figure 2: Highly correlated topics are colored by the same color respectively.

look up the most similar words representing each topic, similar to the way LDA (Blei et al., 2003) represents topics as word distributions. (Kirchhoff, 2019)

#### 4 The SocialVisTUM Toolkit

**Visualization** Figure 2 shows an example of the visualization and labeling tool. Topics are represented as nodes with according labels, and the number of texts assigned to the topics is given in parentheses next to the label. The node size increases based on the number of *topic occurrences*. The edges of the topic connections are labeled by the *topic correlations*. The link thickness increases with a higher positive correlation. A graph layout based on repelling forces between nodes helps to avoid overlaps, which is especially useful when many nodes and links are displayed. A second force keeps the graph centered. Users can also move nodes around to get a more comprehensible overview. (Kirchhoff, 2019)

**Topic Nodes and Correlations** After training the ABAE model, the sentences are assigned to topics based on the maximum topic probability from  $\mathbf{p}_t$ , see formula 2. The correlation between two topics  $i$  and  $j$  is calculated based on the probabilities  $(\mathbf{p}_t)_i$  and  $(\mathbf{p}_t)_j$  of each given sentence  $t$ . This yields a value in the range of  $[-1, 1]$  for every pair of topics specifying the strength of the corresponding relatedness. (Kirchhoff, 2019)

**Hiding Insignificant Topics** An occurrence threshold slider defines the percentage of sentences that must be about a topic to display the associated node. Another slider can be used to set the correlation threshold to define the required positive or/and negative correlation to display the associated connections. These sliders are especially helpful to maintain a clear visualization by limiting the number of shown topics and connections when there

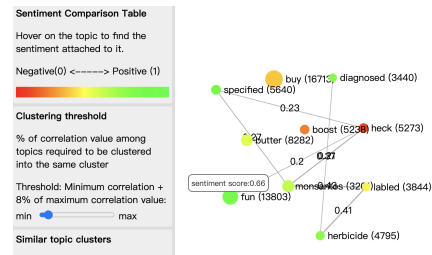


Figure 3: The sentiment for each topic is shown as green (positive) and red (negative).

are many of them available.

**Topic Inspection** Users can double-click a node to receive additional information about a topic, i.e., representative words and sentences, as shown in Figure 5 on the left and right respectively. As representative words, the top 10 words are shown sorted by the distance of their embeddings to the selected topic centroid in ascending order. The representative sentences are the ones with the highest probability from  $p_t$  for the given topic. During topic inspection mode, only nodes that are connected to the clicked node and the associated links are displayed. A double click on the same node brings back the whole graph again.

**Colorization of Topic Nodes** In an updated version of SocialVisTUM, we introduce two meaningful colorings of the topic nodes for correlated topic clustering and sentiment analysis.

Firstly, we perform a hierarchical clustering algorithm such that those topics which are strongly correlated with each other are colored in one and the same color respectively. A dynamic slider GUI element helps to adjust the correlation threshold accordingly. One example outcome is shown in 2.

Secondly, we perform sentiment analysis using the Valance Aware Dictionary and sEntiment Reasoner or VADER method (Gilbert, 2014). It is a lexicon and rule-based sentiment analysis tool that is specifically attuned to sentiments expressed in social media settings. It gives positive, negative, and neutral scores to represent the proportion of text in that sentence that falls in these categories. For each sentence, we use the compound score, i.e., the sum of all lexicon ratings normalized between -1 (most negative) and +1 (most positive). We then calculate the average sentiment score for each topic based on all respective topic sentences. In Figure 3, positive sentiment is shown as green topic nodes, and negative as red. (Roy and Zhao, 2020)

Topic Label	Representative words	# Hypernyms
animal (102)	insect, ant, habitat, rodent, herbivore	218
compound (91)	amino, enzyme, metabolism, potassium, molecule	158
chemical (74)	fungicide, insecticide, weedkiller, preservative, bpa	131
systematically (0)	systematically, adequately, cleaned, properly, milked	0

Table 1: Example topics, automatically assigned topic labels, and representative words. The value next to the topic label denotes how often the label occurs as a shared hypernym. The number of hypernyms on the right tells in how many word comparisons a shared hypernym is identified. Taken from (Kirchhoff, 2019).

**Automatic Topic Labels** We introduce an approach to label topic nodes automatically. It is based on *shared hypernyms*, i.e., the lowest common denominator for words, which we identify using the representative topic words denoted in the previous paragraph and the lexical database WordNet (Miller, 1995). First, we retrieve the hypernym hierarchy for every representative topic word, as shown in Figure 4, and compare every word with every other word in the word list. Next, at each comparison, we save the hypernym with the lowest distance to the compared words in the hypernym hierarchy. We denote these as *shared hypernyms*. We only consider hypernyms if their distance to the word is smaller than half of the distance to the root hypernym to avoid unspecific labels like *entity* and *abstraction*. Eventually, we employ the hypernym that occurs most often as topic labels. If no hypernym can be identified, we use the most representative word instead. In the example shown in figure 4, we identify *dairy\_product* as the lowest shared hypernym of *yoghurt* and *butter*, and *food* as lowest shared hypernym of *yoghurt* and *bread*. (Kirchhoff, 2019)

The quality of a shared hypernym chosen as topic label can be approximated by inspecting the number of its hypernym occurrences – see table 1. Topic labels occurring frequently as shared hypernym are usually suitable (e.g., animal (102) and compound (91)) in contrast to topic labels that rarely occur (e.g. group\_action (9) or smuckers (0)). Thus, we conclude that the number of hypernym occurrences of each topic is suitable to estimate the topic coherence for hyperparameter optimization – see section 5.2 later on. (Kirchhoff, 2019)

**Changing Labels** To change the label of a topic, the user can click on the associated label of a node. This opens a prompt allowing the user to change the topic label. The user can download a JSON file with the updated labels by clicking on the *Create*

*file* button on the sidebar. (Kirchhoff, 2019)

## 5 Case Study

We demonstrate SocialVisTUM’s potential for social media data exploration on a new data set.

### 5.1 Data Set

We crawled online user comments on organic food from multiple forums (e.g., Reddit, Quora, Disqus) and the comment sections of news websites (The Washington Post, The New York Times, Chicago Tribune, HuffPost, and many more). The goal is to discover the discussed topics and opinions in social media regarding the organic food consumption.

Relevant articles from the platforms are found by the search terms ”organic food”, ”organic agriculture”, and ”organic farming”. We further filtered for domain relevance by applying naive Bayes classification on bag of words trained on 1000 random and accordingly labeled texts (84.70% accuracy with 10-fold cross validation). From the left texts, we retain comments containing either of the words *food* and *organic*. The left data set consists of 515.347 sentences totaling 83.938 posts, which are used to train the ABAE topic model. We use the 300-dimensional pre-trained GloVe embeddings and fine-tuned them on the data. (Kirchhoff, 2019)

### 5.2 Hyperparameter Estimation

Some hyperparameters of the utilized ABAE topic model are the number of topic clusters and the vocabulary size. To optimize these automatically, we define a new metric, the *average number of shared hypernyms* (ANH). We first derive the frequency of all shared hypernyms for each topic, which is already done for automatic topic labeling in section 4. The ANH is the sum of hypernym frequencies over all topics divided by the number of topics. (Kirchhoff, 2019)

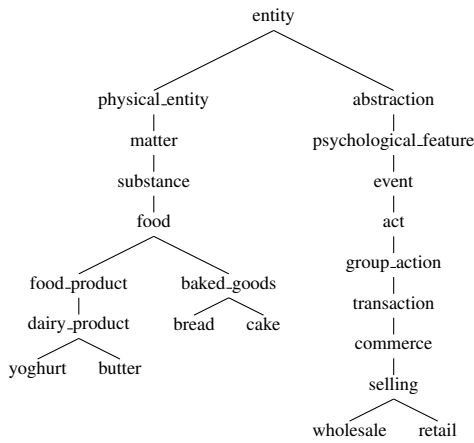


Figure 4: Hypernyms for *yoghurt*, *butter*, *bread*, *cake*, *wholesale*, and *retail*. Taken from. (Kirchhoff, 2019)

In our case study, we identified the following advantages of ANH over the widely used coherence score (CS). First, an increasing number of topics does not always increase the ANH, as a high number of topics leads to many incoherent topics with fewer shared hypernyms, i.e., a lower ANH. Second, a medium-sized vocabulary (~10,000 words) produces the most coherent topics according to ANH and manual inspection. Table 2 shows an excerpt of the results for varying parameters.

### 5.3 Interpretation

We applied SocialVisTUM to our case of organic food yielding the topics displayed in figure 1. A consumer researcher in the domain of organic food manually refined the automatic labeling based on the most similar words of each topic. The topics reflect previous findings of a qualitative content analysis on a small sub-sample of our data set (Danner and Menapace, 2020). The correlated topics allow market researchers to investigate the context in which topics are discussed.

Figure 5 takes a closer look at the example topic *pesticides*, which is concerned with different pesticides and their toxicity. The topic *pesticides* is correlated with the topic *organic\_production\_standards*, which references different organic or related production methods, such as bio-dynamic, hydroponic, or bio-intensive agriculture. This correlation suggests that, for the commenting users in our data set, the non-use of chemical-synthetic pesticides is an important characteristic of organic compared to non-organic production. Further topics correlated with *pesticides* propose that the commenting users are concerned

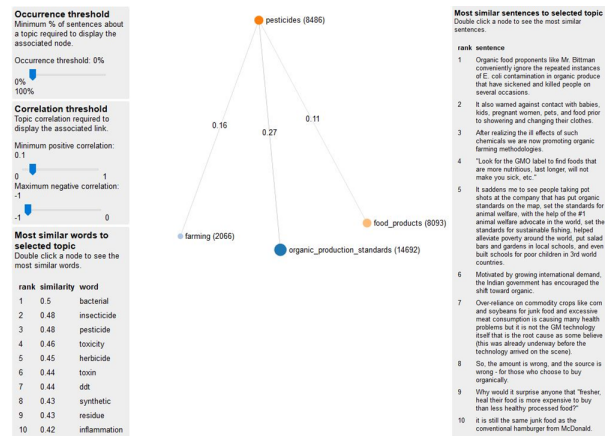


Figure 5: SocialVisTUM applied to our organic food use case. Topic inspection of the *pesticides* topic.

about the use of *pesticides* in *farming* and that they discuss the issue of *pesticides*, possibly the residues thereof, in the context of different *food\_products*.

## 6 Conclusion

In this paper, a case of the proposed SocialVisTUM demonstrates the visualization of coherent topics on a given corpus of social media texts about organic food. The graph-based visualization with topics as nodes and topic correlations as edges reflects the topics and patterns found in a related qualitative content analysis (Danner and Menapace, 2020). The presentation of additional topic information, such as word lists, representative sentences, topic importance, and meaningful predefined labels, provide a basis for the understanding and interpretation of a topic for domain experts. The integrated hyperparameter optimization automatically yields interpretable topics and helps tailoring the model to the given data set. For future work, we plan to evaluate the correlated topics on other corpora and in other use cases. In addition to Pearson correlation, other correlations could improve the approach. We plan to integrate multi-lingual word features, such as BERT (Devlin et al., 2018), for cross-cultural comparisons.

## Acknowledgments

We thank inovex GmbH for supporting the research of this paper by providing computational resources, Robert Pesch and Martin Kirchhoff for their important contributions. Paragraphs adopted from student works are followed by the corresponding citation after the punctuation.

## References

- Stefanos Angelidis and Mirella Lapata. 2018. [Summarizing opinions: Aspect extraction meets sentiment prediction and they are both weakly supervised](#). *CoRR*, abs/1808.08858.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Jonah Berger, Ashlee Humphreys, Stephan Ludwig, Wendy W. Moe, Oded Netzer, and David A. Schweidel. 2020. [Uniting the tribes: Using text for marketing insight](#). *Journal of Marketing*, 84(1):1–25.
- D. Blei and J. Lafferty. 2006. [Correlated topic models](#). In *Advances in Neural Information Processing Systems 18*, volume 18, page 147, Cambridge, MA: MIT; 1998.
- D. Blei and J. Lafferty. 2007. A correlated topic model of science. *Annals of Applied Statistics*, 1:17–35.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.
- Gerlof Bouma. 2009. Normalized (pointwise) mutual information in collocation extraction. *Proceedings of GSCL*, pages 31–40.
- Zhiyuan Chen, Arjun Mukherjee, and Bing Liu. 2014. Aspect extraction with automated prior knowledge learning. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 347–358.
- Hannah Danner, Gerhard Hagerer, Florian Kasischke, and Georg Groh. 2020. [Combining content analysis and neural networks to analyze discussion topics in online comments about organic food](#). In *Conference Proceedings of "3rd International Conference on Advanced Research Methods and Analytics"*, Valencia, Spain. Editorial Universitat Politècnica de València.
- Hannah Danner and Luisa Menapace. 2020. Using online comments to explore consumer beliefs regarding organic food in german-speaking countries and the united states. *Food Quality and Preference*, 83:103912.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ran Ding, Ramesh Nallapati, and Bing Xiang. 2018. Coherence-aware neural topic modeling. *arXiv preprint arXiv:1809.02687*.
- Mainak Ghosh. 2020. Multilingual opinion mining on social media comments using unsupervised neural clustering methods. Master’s thesis, Technical University Munich, Department of Informatics, Boltzmannstr. 3, 85748 Garching, Germany. Unpublished; supervised by Gerhard Hagerer and Georg Groh.
- CJ Hutto Eric Gilbert. 2014. [Vader: A parsimonious rule-based model for sentiment analysis of social media text](#). In *Eighth International Conference on Weblogs and Social Media (ICWSM-14)*.
- Junxian He, Ying Huang, Changfeng Liu, Jiaming Shen, Yuting Jia, and Xinbing Wang. 2016. [Text network exploration via heterogeneous web of topics](#). In *ICDM Workshops*, volume abs/1610.00219, pages 99–106. IEEE Computer Society.
- Ruidan He, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. 2017. An unsupervised neural attention model for aspect extraction. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 388–397.
- Giannis Karamanolakis, Daniel Hsu, and Luis Gravano. 2019. [Leveraging just a few keywords for fine-grained aspect detection through weakly supervised co-training](#). *CoRR*, abs/1909.00415.
- Martin Kirchhoff. 2019. Summarizing opinions using neural topic modeling and graph-based visualizations. Master’s thesis, Technical University of Munich, Arcisstraße 21, 11. Unpublished; supervised by Gerhard Hagerer, M.Sc, and Prof. Dr. Georg Groh.
- Wei Li and Andrew McCallum. 2006. [Pachinko allocation: Dag-structured mixture models of topic correlations](#). In *Proceedings of the 23rd International Conference on Machine Learning, ICML '06*, page 577–584, New York, NY, USA. Association for Computing Machinery.
- Shixia Liu, Xiting Wang, Jianfei Chen, Jim Zhu, and Baining Guo. 2014. [Topicpanorama: A full picture of relevant topics](#). In *IEEE VAST*, pages 183–192. IEEE Computer Society.
- Ling Luo, Xiang Ao, Yan Song, Jinyao Li, Xiaopeng Yang, Qing He, and Dong Yu. 2019. [Unsupervised neural aspect extraction with sememes](#). In *IJCAI*, pages 5123–5129. ijcai.org.
- Arun S. Maiya and Robert M. Rolfe. 2014. [Topic similarity networks: Visual analytics for large document sets](#). *CoRR*, abs/1409.7591:364–372.
- George Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.
- David Mimno, Hanna M Wallach, Edmund Talley, Miriam Leenders, and Andrew McCallum. 2011. Optimizing semantic coherence in topic models. In *Proceedings of the conference on empirical methods in natural language processing*, pages 262–272. Association for Computational Linguistics.

Joe Murphy, Michael W Link, Jennifer Hunter Childs, Casey Langer Tesfaye, Elizabeth Dean, Michael Stern, Josh Pasek, Jon Cohen, Mario Callegaro, and Paul Harwood. 2014. Social media in public opinion research: Executive summary of the aapor task force on emerging technologies in public opinion research. *Public Opinion Quarterly*, 78(4):788–794.

David Newman, Jey Han Lau, Karl Grieser, and Timothy Baldwin. 2010. Automatic evaluation of topic coherence. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, HLT '10, page 100–108, USA. Association for Computational Linguistics.

Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.

Archishman Roy and Jiaxi Zhao. 2020. Enhancing socialvistum. Unpublished lab course report at Technical University of Munich; supervised by Gerhard Hagerer, M.Sc., and Prof. Dr. Georg Groh.

Keith Stevens, Philip Kegelmeyer, David Andrzejewski, and David Buttler. 2012. Exploring topic coherence over many models and many topics. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 952–961. Association for Computational Linguistics.

## Appendix A. Comparison of Coherence Metrics

# Topics	Voc. Size	CS	ANH
5	1,000	-1104	28.6
	<b>10,000</b>	-765	<b>68.0</b>
	18,000	-403	5.2
15	1,000	-366	33.3
	<b>10,000</b>	-270	<b>40.0</b>
	18,000	-197	33.8
50	1,000	-110	30.4
	<b>10,000</b>	-70	<b>51.8</b>
	18,000	-54	49.7

Table 2: Comparing two coherence metrics: coherence score (CS) and average number of shared hypernyms (ANH). The advantage of ANH is that it has its global optimum always in the middle as opposed to CS. This property is beneficial for hyperparameter optimization.