

# Social Media, Topic Modeling and Sentiment Analysis in Municipal Decision Support

Miloš Švaňa

VSB - Technical University of Ostrava  
Faculty of Economics  
17. listopadu 2172/15, 708 00 Ostrava, Czechia  
Email: milos.svana@vsb.cz

**Abstract**—Many cities around the world are aspiring to become *smart*. However, smart initiatives often give little weight to the opinions of average citizens.

Social media are one of the most important sources of citizen opinions. This paper presents a prototype of a framework for processing social media posts with municipal decision-making in mind. The framework consists of a sequence of three steps: (1) determining the sentiment polarity of each social media post (2) identifying prevalent topics and mapping these topics to individual posts, and (3) aggregating these two pieces of information into a fuzzy number representing the overall sentiment expressed towards each topic. Optionally, the fuzzy number can be reduced into a tuple of two real numbers indicating the "amount" of positive and negative opinion expressed towards each topic.

The framework is demonstrated on tweets published from Ostrava, Czechia over a period of about two months. This application illustrates how fuzzy numbers represent sentiment in a richer way and capture the diversity of opinions expressed on social media.

## I. INTRODUCTION

**M**ANY CITIES around the world are aspiring to become *smart*. A *Smart City* is characterized by an extensive use of information technologies to support municipal decision-makers in effective resource utilization [17]. This ICT support can be applied in many areas, including traffic and transportation, waste management, accommodation or culture.

Ideally, municipal planning involves multiple stakeholder groups: local authorities, businesses, average citizens and commuters, or environmental activists. In practice, however, decisions are mostly based on an interplay between authorities and businesses offering Smart City technologies. Opinions of average citizens are often given little weight [9, 18].

Some cities use surveys to gather opinions on at least some issues and projects. Although useful, surveys have many disadvantages. Most importantly, by asking only a predetermined set of questions and letting the respondent choose only from a low number of possible answers they limit the respondent's ability to fully express their thoughts and opinions. This prevents decision makers from serendipitously discovering unexpected issues and ideas.

Social media represent an alternative source of citizen opinions. They offer much greater freedom of expression as users can create new content whenever they want and instead

of being limited by a set of questions, they can use free-form text, images, videos or audio. This freedom together with significant content creation velocity might also be a disadvantage. Relatively simple statistical methods for processing survey data are inadequate for social media. Instead, advanced techniques from fields such as machine learning, computer vision, or natural language processing have to be used.

This paper presents a prototype of a framework for extracting information from social media with the aim to support municipal decision-making. The framework combines topic modeling techniques with sentiment analysis. First, the system detects topics discussed on social media at a specific location. Then it evaluates sentiment towards each topic. To capture uncertainty arising from different people having different opinions on the same topic, this sentiment is modelled as a triangular fuzzy number (TFN). The TFN representation, however, might not be understood by people without background knowledge. Therefore a calculation of a "degree of conformity" with fuzzy sets representing the concepts of positive and negative opinion follows. The result can be interpreted as an "amount" of positive/negative opinion expressed towards a specific topic.

The remainder of the paper is organized as follows: In section II I review literature related to topic modeling, sentiment analysis and their applications in the context of Smart Cities. Section III discusses the framework design. Section IV then describes data used in experiments and preprocessing procedures. Demonstration of the framework application on test data can be found in section V. Finally, section VI discusses the limitations of the framework and future research.

## II. LITERATURE REVIEW

### A. Topic Modeling

Topic modeling can be understood as (1) a "statistical technique for revealing the underlying semantic structure in a large collection of documents", or (2) "a technique comes with group of algorithms that reveal, discover and annotate thematic structure in collection of documents" [12].

In practice, topic modeling involves taking a corpora of text documents and discovering various topics discussed in these documents. Topics are usually represented by a set of

most relevant terms [4]. Once topics are identified, a mapping between topics and individual documents can be created. Depending on the method used, one document can be assigned to one or multiple topics.

There is a plethora of topic modeling methods. In [12], the authors distinguish between two categories: (1) algebraic methods, usually based on some form of word-document matrix factorization and (2) statistical methods. On the other hand, [4] provides a chronological overview of the development of topic modeling methods. Other reviews, e.g., [16] then focus on selecting the best method for a given dataset or application.

One of the most popular topic modeling methods is Latent Dirichlet Allocation (LDA). This method is based on the idea that documents in a given corpora are generated by a probabilistic process. Each document can be understood as a mixture of topics, with each topic representing a probability distribution over words from some vocabulary [3]. However, the framework presented in this paper relies on a more recent method called BERTopic [8]. It outperforms LDA on multiple benchmark datasets both in coherence and topic diversity – two common topic modeling evaluation metrics [4].

### B. Sentiment analysis

Sentiment analysis is "an approach that uses Natural Language Processing (NLP) to extract, convert and interpret opinion from a text and classify them into positive, negative or natural sentiment" [7]. There are two main groups of sentiment analysis methods (1) lexicon-based and (2) methods based on supervised learning models.

Lexicon-based methods use sentiment lexicons containing information about sentiment polarity of different words. Individual word polarities in a given document are looked up and then aggregated into an overall document polarity.

Methods based on supervised learning models rely on various machine learning techniques to train a classification model for predicting document sentiment. Naive Bayes is one of the most popular approaches in this category. Supervised methods require a training dataset that contains a ground truth sentiment for each document. Many product and service review websites let their users combine textual reviews with some sort of a numerical rating scale, e.g., a 5-star rating system. This data can be easily used to train a sentiment classifier.

One of the contributions of this paper lies in the application of fuzzy modelling methods in sentiment analysis. Several researchers have already explored this path. In [11] authors describe a fuzzy expert system for sentiment analysis. There have also been attempts at designing a sentiment analysis method that utilizes a fuzzy thesaurus [10]. However, to our best knowledge, fuzzy methods have not been used to aggregate sentiment expressed towards multiple documents.

There have been multiple applications of sentiment analysis in the context of Smart Cities and urban planning [5, 14]. There are many opportunities and challenges related to the use of sentiment analysis in urban planning, including visualiza-

tion, multilingual audiovisual opinion mining, or peer-to-peer opinion mining tools for citizens [2].

## III. FRAMEWORK DESIGN

### A. Topic Modeling

As mentioned, the framework prototype uses a topic modeling method called BERTopic. Originally proposed in [8], BERTopic solves the problem of topic modeling by combining word embeddings with hierarchical clustering. The procedure consists of the following steps:

- 1) **Create an embedding for each document.** One of the downsides of traditional methods such as LDA is that they represent documents in a bag-of-words fashion. This representation ignores both order of words in a document and their semantic relationships. BERTopic relies on an embedding representation. Embeddings are vectors that are able to somewhat capture the semantics of words or documents. Documents with similar meaning should be represented by similar vectors. As the name suggests, BERTopic uses embeddings based on BERT [6].
- 2) **Reduce embedding dimensionality.** Embeddings can have hundreds or even thousands of dimensions. When it comes to clustering, high dimensionality causes multiple issues. [8]. First, the difference between the distance of the nearest point to a cluster centre and the distance between the furthest point from a cluster centre shrinks [1]. Second, lower number of dimensions leads to better performance both in terms of time and clustering accuracy. BERTopic therefore uses UMAP [15] to reduce the number of dimensions.
- 3) **Use a clustering algorithm to create document clusters.** BERTopic uses HDBSCAN with document embeddings reduced by UMAP as input. HDBSCAN is a hierarchical algorithm able to create a tree of cluster-subcluster structures. BERTopic user can set the number of clusters/topics to be generated by the algorithm. To provide the desired number of clusters, small similar topics are merged together.
- 4) **Create topic representations.** Similarly to other topic modeling methods, BERTopic represents each topic as a list of words. To find words best describing a given topic BERTopic uses a modified TF-IDF score [8] calculated as:

$$W_{t,c} = tf_{t,c} \log\left(1 + \frac{A}{tf_t}\right) \quad (1)$$

where  $tf_{t,c}$  represents the frequency of term  $t$  in cluster  $c$ .  $A$  is the average number of words per cluster and  $tf_t$  is the total frequency of term  $t$  across all clusters.

### B. Sentiment Analysis

The TextBlob<sup>1</sup> Python library was selected for the framework prototype. This choice was made mainly for pragmatic

<sup>1</sup><https://github.com/sloria/TextBlob>

reasons - ease of installation, use and integration with other parts of the framework.

TextBlob offers two sentiment analysis methods: a lexicon and pattern-based one and a pretrained Naive Bayes model. In the former, TextBlob uses a polarity lexicon and structural patterns to determine both the degree of polarity and the degree of subjectivity. The Naive Bayes model was trained to classify movie reviews as positive or negative. Instead of providing degrees of polarity and subjectivity as output, this method returns the probabilities of a given text being positive and negative. Being the default, the lexicon and pattern-based method was also used in presented experiments. The output of this model carries a certain degree of uncertainty which can be exploited when aggregating sentiment analysis results with fuzzy methods.

### C. Fuzzy Aggregation

Once sentiment polarity and topics are extracted from a set of social media posts, both pieces of information are combined to provide an overall view of what is being discussed in a given municipality and whether the population perceives a given topic positively or negatively.

Simple aggregation metrics such as arithmetic mean would lead to a loss of information. For example, there might be a controversial topic with many positive, but also some negative opinions. Arithmetic mean might present the topic sentiment as slightly positive. When presented with this information, the user cannot tell whether the aggregated opinion is slightly positive because a majority of individual opinions is slightly positive or because there are many positive opinions counterbalanced by a small number of negative opinions.

To address this issue, topic sentiment is modeled as a triangular fuzzy number (TFN). Assuming we know both the sentiment polarity of each social media post, and their topic distribution, the TFN core can be determined as a weighted mean with topic distribution percentages serving as weights. For instance, consider 3 topics with polarities and topic distributions depicted in table I. The core of the TFN representing the sentiment towards topic 1 can be calculated as:

$$m_{t1} = \frac{0.5 \cdot 0.5 + 0.35 \cdot 0.3 - 0.2 \cdot 0.2}{0.5 + 0.3 + 0.2} = 0.315 \quad (2)$$

Weighted standard deviation is then used to determine the TFN support interval. As in the case of determining the TFN core, the degree to which a given post belongs to a given topic should determine the strength of its influence on the shape of the support interval. Therefore a weighted variant of standard deviation is used, the general formula of which being:

$$\sigma = \sqrt{\frac{\sum_{i=1}^N w_i (x_i - x^*)^2}{\frac{M-1}{M} \sum_{i=1}^N w_i}} \quad (3)$$

with  $N$  representing the total sample size,  $x^*$  representing weighted mean, and  $M$  the number of non-zero weights.

Once the weighed standard deviation is known, the TFN support interval is calculated as:

$$[m_{t1} - s\sigma; m_{t1} + s\sigma] \quad (4)$$

where  $s$  is a positive real number that determines the scaling between the length of the support interval and weighted standard deviation. Its influence on the aggregation result is a subject of further research.

TABLE I  
EXAMPLE OF AN INPUT FOR THE AGGREGATION PROCEDURE

	Polarity	Topic distribution		
		Topic 1	Topic 2	Topic 3
Post 1	0.5	0.5	0.3	0.2
Post 2	0.35	0.3	0.4	0.3
Post 3	-0.2	0.2	0.2	0.6

Taking inspiration from [19], the next step of the aggregation procedure is the calculation of a degree of conformity between the TFN describing topic sentiment and fuzzy sets representing the concepts of a positive and negative opinion. The result of this step is a tuple of two numbers from interval  $[0; 1]$  that could be interpreted as the overall "amount" of positive and negative opinion expressed towards a given topic. Although there is some loss of information when compared to the original TFN, the two numbers are more likely to be understood by a layperson with no knowledge of fuzzy set theory.

The fuzzy metric of possibility [13] is used to calculate the degree of conformity, e.g.:

$$Pos(\tilde{A}, \tilde{P}O) = \sup_{x \in X} \min(\mu_{\tilde{A}}(x), \mu_{\tilde{P}O}(x)) \quad (5)$$

$A$  denotes aggregated topic sentiment with  $\mu_{\tilde{A}}(x)$  being its membership function. Similarly  $PO$  represents the concept of a positive opinion with a membership function  $\mu_{\tilde{P}O}(x)$ .

## IV. DATA AND PREPROCESSING

Social media posts published on Twitter were used to demonstrate the proposed framework. This social network was chosen specifically because it lets researchers easily access its data through an API. Using this API, a dataset of approximately 3000 tweets published by users located in Ostrava, Czechia was created. These tweets cover the period of January and February of 2023.

Given the location, most tweets in the dataset are written in Czech. There are, however, several tweets written in other languages, such as English, Slovak or Polish. To address the issue of multilingualism, as well as the fact that TextBlob uses an English lexicon when analyzing sentiment, the DeepL Translator<sup>2</sup> API was used to automatically translate each post to English.

Several preprocessing steps were then applied on translated tweets:

<sup>2</sup><https://www.deepl.com/translator>

- removal of tweets that are shorter than 60 characters under the assumption that it might be difficult to reliably extract sentiment and topics from short tweets,
- removal of non-letter characters,
- removal of hashtags, usernames and URLs,
- turning all text into lowercase,
- removal of stop words.

## V. FRAMEWORK DEMONSTRATION

The framework prototype was implemented as a Jupyter notebook. To summarize, after the input data was pre-processed, the following sequence of steps was applied:

- 1) TextBlob was used to determine the sentiment polarity of each post. Polarity of each post is represented as a real number from a  $[-1, 1]$  interval with -1 representing absolutely negative polarity and positive number representing absolutely positive polarity.
- 2) BERTopic was used to identify topics and create a mapping between topics and individual tweets. Provided output follows the format in table I.
- 3) A set of TFNs representing aggregated topic sentiment was constructed. The scaling constant  $s$  in equation 4 was set to 1.
- 4) Degrees of conformity between each topic's TFN and the concepts "positive opinion" and "negative opinion" were determined.

Overall, BERTopic identified 42 different topics. Topics with the highest overall "mass" calculated as the sum of the topic's probability across all tweets are listed in table II.

Presented topic distribution intuitively makes sense. During given time period, one of the most dominant topics in public discourse was the upcoming presidential election which corresponds to the first topic in table II. One can also see that not all topics are necessarily relevant to municipal decision making, for example topic no. 3 in the table. At the same time, topics such as no. 9 are difficult to interpret.

TFNs representing the several aggregated topic sentiments are depicted in figure V. It can be deduced that the topic *vote, election, politics, party* is perceived quite positively and compared to topic *pay, wage, pension, live* it also has a narrower support. This indicates a lower level of opinion diversity. The latter topic is also perceived most negatively, at least among the topics displayed in the figure.

Table II also contains information about the conformity with the concepts of positive and negative opinion. The membership function of the fuzzy set representing positive opinion has a value of 0 until polarity value 0, then grows linearly to 1 until polarity reaches 0.2 and then has a value of 1. The "negative opinion" fuzzy set is its mirror image.

It can be seen that the information provided by these numbers tells a story similar to the TFN visualization. For example, when comparing topics *school, education, teacher, class* and *area, city, building, ostrava*, one could conclude that the former is not only perceived more positively, but that it is also less controversial given the lower "amount" of negative opinion.

## VI. DISCUSSION

As demonstrated, the framework prototype can be used to create a representation of opinions expressed towards different topics on social media posted by citizens of a specific municipality. The prototype, however, employs certain simplifications that should be addressed in future research.

The TFNs representing topic sentiments are currently symmetric. This might not be the best reflection of reality. Ways of making the TFNs asymmetric should be explored. Metrics such as weighted skewness or semivariance could be used. Next, Twitter users have an option to "like" or "retweet" a post created by someone else. These actions can represent approval and could be therefore used as aggregation weights. Finally, other topic-modeling and sentiment-analysis methods should be tested and compared to the existing TextBlob+BERTopic stack. Moreover, TextBlob provides additional information beyond sentiment polarity: the subjectivity degree of each document. This metric could again be used as weight.

Framework output can be used directly by municipal decision makers to make more informed decisions. However, there are other possible applications. One of them is comparison of different municipalities. It should be possible to compare the TFNs of a specific topic across multiple cities. If additional information such as municipal budgets is available, one could also deploy methods such as Data Envelopment Analysis to evaluate their efficiency. Topic sentiments could be also aggregated into an overall sentiment expressed towards everything happening in a given municipality.

However, using the proposed framework as a replacement for other methods of gathering citizen opinion might not be the best course of action. Instead, the framework should play a complementary role, as social media users might not accurately represent the overall population. Groups such as the elderly might be underrepresented. At the same time, methods such as surveys can provide biased information too. Combining the proposed framework with surveys might lead to a better overall representation of citizen opinion than either of these methods separately.

## VII. ACKNOWLEDGMENT

This paper was supported by the SGS project No. SP2023/078. This support is gratefully acknowledged.

## REFERENCES

- [1] C. C. Aggarwal, A. Hinneburg, and D. A. Keim. On the surprising behavior of distance metrics in high dimensional space. In J. Van den Bussche and V. Vianu, editors, *Database Theory — ICDT 2001*, pages 420–434, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.
- [2] K. B. Ahmed, A. Radenski, M. Bouhorma, and M. B. Ahmed. Sentiment analysis for smart cities : State of the art and opportunities. In *Int'l Conf. Internet Computing and Internet of Things (ICOMP'16)*, 2016.
- [3] D. M. Blei, A. Ng, and M. I. Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022, 2001.

TABLE II  
MOST PREVALENT TOPICS IDENTIFIED BY BERTOPIC IN TWEETS PUBLISHED FROM OSTRAVA.

No.	Topic	Prevalence	Positivity	Negativity
1	'vote', 'election', 'politics', 'party', 'ano', 'government', 'people'	41.98	0.997	0.190
2	'czech', 'republic', 'czech republic', 'czechs', 'czk', 'inflation'	33.68	0.810	0.434
3	'boy', 'date', 'girl', 'crying', 'apologize', 'shes', 'ive'	27.19	0.812	0.390
4	'school', 'education', 'teacher', 'class', 'educated', 'academics', 'students'	25.93	0.897	0.287
5	'game', 'field', 'leader', 'go', 'games', 'team', 'match'	25.89	0.839	0.390
6	'area', 'city', 'building', 'ostrava', 'construction', 'also', 'mappa'	25.32	0.745	0.473
7	'hydrogen', 'gt', 'car', 'anymore', 'hydrogen bike', 'product', 'hydrogen technology'	19.47	0.741	0.459
8	'something nothing', 'number', 'nothing', 'something', 'old figure', 'number thumbs'	18.81	0.760	0.272
9	'miracles', 'ribbons', 'expect miracles', 'political responsibility', 'imagine'	18.72	0.757	0.344

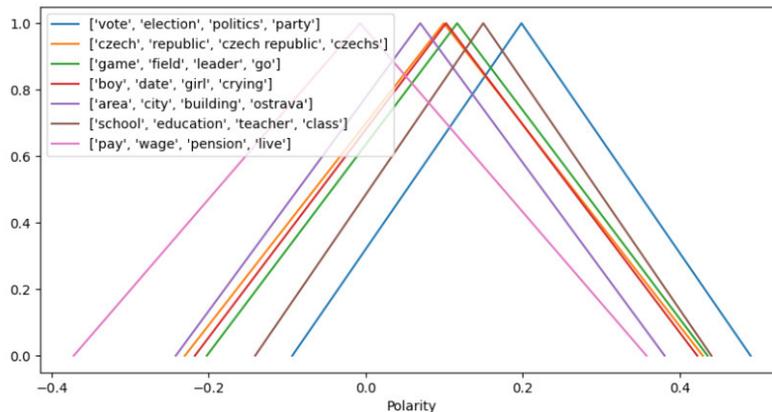


Fig. 1. Depiction of TFNs representing overall sentiment towards selected topics extracted from tweets published from Ostrava

- [4] R. Churchill and L. Singh. The Evolution of Topic Modeling. *ACM Computing Surveys*, 54(10s):1–35, Jan. 2022.
- [5] M. Dahbi, R. Saadane, and S. Mbarki. Social media sentiment monitoring in smart cities: an application to Moroccan dialects. In *Proceedings of the 4th International Conference on Smart City Applications*, pages 1–6, Casablanca Morocco, Oct. 2019. ACM.
- [6] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding, 2019.
- [7] Z. Drus and H. Khalid. Sentiment Analysis in Social Media and Its Application: Systematic Literature Review. *Procedia Computer Science*, 161:707–714, 2019.
- [8] M. Grootendorst. Bertopic: Neural topic modeling with a class-based tf-idf procedure, 2022.
- [9] G. R. Halegoua. *Smart cities*. The MIT Press, 2020.
- [10] H. M. Ismail, B. Belkhouche, and N. Zaki. Semantic Twitter sentiment analysis based on a fuzzy thesaurus. *Soft Computing*, 22(18):6011–6024, Sept. 2018.
- [11] C. Jefferson, H. Liu, and M. Cocea. Fuzzy approach for sentiment analysis. In *2017 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, pages 1–6, Naples, Italy, July 2017. IEEE.
- [12] P. Kherwa and P. Bansal. Topic Modeling: A Comprehensive Review. *ICST Transactions on Scalable Information Systems*, 0(0):159623, July 2018.
- [13] G. Klir and B. Yuan. *Fuzzy Sets and Fuzzy Logic: Theory and Applications*. Prentice Hall PTR, 1995.
- [14] M. Li, E. Ch'ng, A. Chong, and S. See. The new eye of smart city: Novel citizen Sentiment Analysis in Twitter. In *2016 International Conference on Audio, Language and Image Processing (ICALIP)*, pages 557–562, Shanghai, China, July 2016. IEEE.
- [15] L. McInnes, J. Healy, and J. Melville. Umap: Uniform manifold approximation and projection for dimension reduction, 2020.
- [16] I. Vayansky and S. A. Kumar. A review of topic modeling methods. *Information Systems*, 94:101582, Dec. 2020.
- [17] T. M. Vinod Kumar, editor. *Smart Economy in Smart Cities: International Collaborative Research: Ottawa, St.Louis, Stuttgart, Bologna, Cape Town, Nairobi, Dakar, Lagos, New Delhi, Varanasi, Vijayawada, Kozhikode, Hong Kong*. Advances in 21st Century Human Settlements. Springer Singapore, Singapore, 2017.
- [18] K. S. Willis. Whose Right to the Smart City? In P. Cardullo, C. Di Felicianantonio, and R. Kitchin, editors, *The Right to the Smart City*, pages 27–41. Emerald Publishing Limited, June 2019.
- [19] F. Zapletal, M. Hudec, M. Švaňa, and R. Němec. Three-level model for opinion aggregation under hesitance. *Soft Computing*, Feb. 2023.