Renáta Németh* & Domonkos Sik**

Beyond 'latent thematic structure': Expanding the interpretation of the topic model towards pragmatics

Intersections. EEJSP
10(4): 148–159.
https://doi.org/10.17356/ieejsp.v10i4.1323
https://intersections.tk.hu

* [nemeth.renata@tatk.elte.hu] (Eötvös Loránd University, Faculty of Social Sciences, ELTE Research Center for Computational Social Science)
** [sik.domonkos@tatk.elte.hu] (Eötvös Loránd University, Faculty of Social Sciences, ELTE Research Center for Computational Social Science)

Abstract

According to the textbook definition, a topic model aims to uncover the underlying topics of a corpus. Despite its widespread use across disciplines, the nature of these 'topics' has remained relatively underdefined. This research note attempts to fill this gap, drawing on empirical evidence to elucidate the practical application of the model. We argue that the frequency of terms within texts is influenced not only by their theme but also by factors such as genre and context, thus extending the notion of 'latent topics' beyond referential-semantic boundaries to include pragmatic considerations. Through case studies focusing on different genres, such as parliamentary speeches and online forums, we demonstrate the importance of pragmatics, which is often overlooked in well-known early applications that deal predominantly with formal written texts such as newspaper articles or academic papers.

Keywords: natural language processing, topic model, model interpretation, pragmatics

1 Problem statement

One of the best-known natural language processing (NLP) models, the topic model (Blei et al., 2003; Lafferty & Blei, 2009), aims to define the latent 'thematic structure' of a corpus of texts, i.e., to identify latent 'topics' based on the terms that occur in the texts. Many variants of the model have been developed over the last two decades and are widely used in a wide range of disciplines, but few papers have addressed the question of what is meant by 'thematic structure' and what exactly are the 'topics' that are identified. According to a recent systematic review, most researchers use topic models in a suboptimal manner due to the lack of methodologically elaborated interpretative protocols (Laureate et al., 2023). According to more profound criticism, topic modeling in its current form is usually based on 'unrealistic assumptions,' cannot be validated, does not model themes or content, and is uncontrollably affected by apophenia and confirmation bias (Shadrova, 2021).

While not so dismissive, we tend to agree with the direction of these criticisms: using topic models in social scientific research requires a lot of methodological and epistemological reflection to clarify the capabilities and limitations of the method. The present paper aims to contribute to these reflections from the perspective of researchers who use topic models for analyzing various textual corpora. While working on various corpora, we have realized both the interpretative and epistemological uncertainties of topic models and attempted to find local solutions to emerging problems. In our experience, complementing quantitative results with iterations of qualitative analysis can help to decrease the above-mentioned uncertainties.

The purpose of this short research note is to explore these issues based on our empirical experience to support the practical use of the model. After presenting the views of some authors on this issue, we will argue that the higher/lower frequency of certain terms within texts depends not only on the narrowly defined topic/theme of the text but also on its form, genre, and context; 'latent topics' may therefore differ not only in referential-semantic terms (which is the narrow sense of 'topic,') but also in pragmatic terms. In the case studies to which we refer, we will see examples of both the role of genre (type of parliamentary speeches), the importance of word forms (the characteristically different roles of nouns and verbs), and, more generally, the usability of speech act theory in this context.

2 A heuristic description of the topic model

According to its authors (Blei et al., 2003; Lafferty & Blei, 2009), the topic model aims at identifying latent themes (in NLP terms, 'topics') in a corpus. Statistically, 'latent topics' are probability distributions over terms in the dictionary. The model attempts to maximize the distinctiveness of topics, i.e., it tries to distribute the words across topics in such a way that the topics are most clearly distinct from each other. A given topic can be identified by the terms that are most specific (most relevant) to it. For example, the co-occurrence of words such as water, farmland, plan, GMO, floods, and hectare in a subset of parliamentary speeches indicates that an agricultural topic is being discussed (see Németh et al., 2025). Thus, the co-occurrence of terms is important in identifying topics. Because of this property, the model stands in sharp contrast to dictionary-based approaches that examine the frequency of predefined dictionary items in texts; for example, political scientists' populist dictionaries used to identify populist politicians.

The topic model represents text with a simple 'bag of words' model that does not take into account word order or sentence tagging but maintains multiplicity, i.e., it keeps track of the number of times each term occurs in the whole text. Therefore, the model requires preprocessing of the text, such as the removal of punctuation and inflectional endings (stemming or lemmatization) as well as common words like conjunctions and articles (stopword removal), and the combining of frequently co-occurring terms and multi-element proper nouns into single terms (significant n-gram identification and named entity recognition).

One of the assumptions of the model is that there is a finite number of topics that 'generate' the texts. Texts can be identified as mixtures of topics, with as few topics as possible for a given text and as few characteristic terms as possible for a given topic. By estimating the parameters of the model, it is possible to find these topics (and the most relevant terms associated with them, i.e., words that are characteristic of the topics) and to estimate the extent to which a given text is associated with each topic (this is the topic contribution indicator). The contents (labels) of the topics estimated by the model are then assigned to each topic by the researcher by interpreting the most relevant terms in the topics and by qualitatively processing the texts that are most representative of the topics in terms of their topic contribution. The relevance of a term (Sievert & Shirley, 2014) is defined as the sum of the topic-specific frequency of a term and a penalty term, which is an increasing function of the overall frequency of the term. Obviously, terms that occur frequently in a given topic but are also very frequent in the corpus as a whole are less relevant to the topic. Ideally, the interpretation is straightforward for the researcher who is familiar with the field, and the topics speak for themselves.

The number of topics is an input parameter of the model. One of the main aspects of fitting and interpreting topic models is, therefore, determining the 'correct' or rather 'optimal' number of topics. This decision can be justified objectively based on the number of clusters in cluster analysis. Choosing too low a number may lead to overly general topics, while choosing too high a number may lead to fragmented, redundant topics. Statistical metrics may also be used to support the decision, such as perplexity, which tests how successfully the model predicts new data, or various coherence metrics, which try to quantify the semantic similarity between the most relevant words in a given topic. The optimal number of topics is then determined by the maximum (or rather the local maximum) of the metrics computed for models with different numbers of topics. However, a mixed approach that also uses a qualitative assessment of the interpretability of the topics often yields better results when the qualitative assessment of the interpretability of the topics is also applied.

There are several different versions of the topic model, which differ in their statistical assumptions; historically, the first and perhaps most commonly used is Latent Dirichlet Allocation (LDA; Blei et al., 2003). More recent variants allow for the analysis of correlation between topics (correlated topic model) or variation over time (dynamic topic model) or different distributions by author (author topic model) or the impact of metadata on topic content (structural topic models), etc.

Exceptionally, in the field of text mining, this model has become popular not only in business but also in academia, namely, in humanities applications. Blei and his co-authors already recommended the method to humanists, and indeed, applications in this field appeared very early on – for example, historians have used it for historical journal analysis or diary analysis, and literary scholars for the analysis of poetic texts. The *Journal of Digital Humanities* devoted a special issue to the method as early as 2012.

To illustrate the productivity of the method, we will present a few applications that are characteristically different in terms of both subject and corpus. Mützel, in the Berlin project (2015), analyzed restaurant reviews of the city's restaurants going back 20 years. The LDA topic model helped to find changes in trends over a longer time span and also allowed the researcher to turn to the corpus qualitatively (through traditional reading) to find a turning point. The model revealed a dominant dynamic, namely that people had begun to pay more attention to the 'new German cuisine' than to local quality and atmosphere with a focus on regional and seasonal ingredients. Light and Cunningham (2016) looked at speeches given at the Nobel Peace Prize acceptance ceremony, which they argue are increasingly associated with globalization and neoliberalism, shifting from earlier (e.g., Christian) schemas. Blevins (2010) applied the method to the diary of Martha Ballard, who lived in the 18th century. The diary contains almost 10,000 entries, so for a systematic overview, it is really helpful to have access to automated tools. The topics separated by the topic model were labeled with terms such as death, housework, and emotions, and it was seen, for example, that entries related to emotions became more frequent as the diary writer aged. An application in media studies is the work of Jacobi et al. (2016), who, from an analysis of *New York Times* articles on nuclear technology from 1945 to 2016, found that LDA is a very useful tool for the relatively fast content review of huge digital text corpora. Political scientist Grimmer (2010) examined how senators explain their work in Washington to voters by simultaneously examining the themes of media texts and senators' press releases. One of Grimmer's findings, which contradicts the theory of some political scientists, is that senators from the same state emphasize similar priorities in their press releases as senators from different states, and the author explains this by suggesting that senators from the same state may rely on similar constituencies. Finally, we present two articles that analyze speeches from the Hungarian Parliament (Németh et al., 2025, focusing on the parliamentary discourse of the Carpathian Basin; and Sik et al., 2024, analyzing the big picture of two decades of Parliament).

3 Previous considerations on the nature of the topics

The following brief scientific history of the introduction and diffusion of the topic model focuses on showing how the first, narrow semantic understanding of the concept of 'topic' necessarily became broader as a result of applications covering an ever wider spectrum and how it became necessary to include the dimension of pragmatics alongside the dimension of semantics in the definition of 'topic.'

In their first paper, the authors of the model take a technical approach to the nature of topics without mentioning any other approach to interpretation: 'We refer to the latent multinomial variables in the LDA model as topics […] so as to exploit text-oriented intuitions, but we make no epistemological claims regarding these latent variables beyond their utility in representing probability distributions on sets of words' (Blei et al., 2003, p. 996).

Later, in a co-authored paper (Lafferty & Blei, 2009), Blei presented the topic model variants using the example of the JSTOR scholarly journal archive, where topics are actually disciplines or sub-disciplines. Here, the presented application of the model operates within a narrow genre field (scientific articles), using a purely semantic approach without the need for a pragmatic dimension. It uses content as a synonym for topic ('To develop the necessary tools for exploring and browsing modern digital libraries, we require automated methods of organizing, managing, and delivering their contents,' p. 71) and refers to latent content structures ('finding useful structure in an otherwise unstructured collection,' p. 71; 'the hidden variables represent the latent topical structure,' p. 73).

Blei's later paper with sociologists is not only a methodological novelty but also an important work in cultural sociology, providing general research experience about the model from a less technical, more interpretation-validation-oriented approach. It is not only about latent themes but also about information reduction, and the importance of interpretation is emphasized ('How can we capture the information we need, reduce its complexity, and provide interpretations that are substantively plausible and statistically validated?', DiMaggio et al., 2013, p. 570). New to the article are the concepts of framing and context, which go beyond the earlier technical semantics of topics ('In applications to the study of culture, substantive interpretability is crucial. Many topics may be viewed as frames (semantic contexts that prime particular associations or interpretations of a phenomenon in a reader) and employed accordingly', p. 578).

To our knowledge, this paper is the first to introduce the concept of relationality in meaning, providing an epistemological approach to the previous technical definition of the topic. According to the authors, meaning is relational in the sense that co-occurrences are important in assigning words to topics, and meanings are derived from these relations rather than from the words themselves. The paper cites structuralist linguistics and Saussure in this context, i.e., the original technical-statistical approach of the topic model is given a new, more complex interpretation in relation to concepts from other disciplines.

Schmiedel et al. (2019) relate the meaning of topics to the linguists' concept of distributional semantics (Firth, 1957) because, according to them, semantic similarities between linguistic items can be determined based on their distributional properties within a topic. This approach claims more than the relational model because it also implies that words occurring in the same context have similar meanings. It is worth noting here that word embedding models (Mikolov et al., 2013), which have gained much importance in recent years, also build on the distributional semantic approach to define vector spaces in which the location of a word depends on its corpus environment, i.e., its meaning. But here, the environment is defined using a small distance (only a few words), whereas topic modeling considers co-occurrences over a much wider range, typically a whole text. Therefore, the meaning relationship detected in topic modeling is less definitive than that defined in word embedding, and it is not possible to identify the meaning of words for topic models.

Blei et al. (2003) and others asserted the 'thematic coherence' of topics; however, the concept of 'coherence' also remained ill-defined in this research. However, later attempts were made to quantify coherence, e.g., Röder et al. (2015). Coherence measures always measure the semantic proximity of the relevant terms of a topic, where proximity is determined by an external reference corpus (e.g., Wikipedia). Röder et al. (2015) measured the validity of coherence metrics by human raters, not directly by going back to the original texts (i.e., by reading the texts belonging to a given topic and judging their coherence) but indirectly by assessing the semantic coherence of the relevant terms of the topic.

Gillings and Hardie (2023) and Brookes and McEnery (2019) have pointed out the shortcomings of this approach to the interpretation of topics, which can also be found in early applications (e.g., DiMaggio et al., 2013). They argue that word lists alone are of limited use for interpreting topics, and a traditional reading of the most representative texts is also necessary. This is supported by our own research experience (for a detailed description of our interpretation and validation procedure, see, e.g., Németh et al., 2021), and other users of topic models also often explicitly refer to mixing methods – see, e.g., Jacobs and Tschötschel (2019) or Chakrabarti and Frye (2017).

Shadrova (2021) gives one of the most comprehensive discussions of the epistemological problems of the nature of topics. Focusing on the claims relevant to the present paper, she distinguishes the meaning of 'topic' according to the discipline in which it is used and sees a marked difference between, for example, information science, social sciences, and literary studies. According to Shadrova, information science understands topics as labels that can be used, for example, in the efficient classification of libraries; in the social sciences, a topic is the meaning derived from words that belong to similar semantic categories, while in literary studies, a topic is woven into a historical or societal context. As Shadrova claims, defining and quantifying co-occurrence in the interpretation of topics is a linguistically daunting task, requiring consideration of genre, text type, length, etc.

We will argue below that the role of societal context is also important in social science applications and that it is, therefore, simplistic to focus only on referential-semantic meaning in interpretation. Indeed, as we have seen in the brief literature review presented above, the first applications of the topic model were in the field of information science and were actually for classifying texts of the same genre (e.g., scientific articles). Moving out of this narrow framework – for example, into the field of cultural sociology – necessitated the use of the concepts of context and framing. After all, the fundamental interpretative challenges of hermeneutic social sciences, as elaborated by the founding fathers (e.g., Weber, 1949), are not rendered obsolete by the new analytical tools. The meaningful interpretation of textual data remains the task of a 'hermeneutic circle' (Grondin, 2015), regardless of whether the corpus consists of interviews or large-scale digital texts. To understand topics as semantic clusters, the 'whole' (represented by the key terms) can be understood from the perspective of the constitutive 'parts' (represented by the individual speech acts), and vice versa, the parts can be understood from the perspective of the whole. The iteration of these rounds of interpretation has the potential to reveal the 'thematic coherence' and fully explore the semantic and performative features of the topics. In what follows, we will show how the application of these principles enriches topic modeling: in our understanding, the semantics of topics indeed include genre and other pragmatic features that have important interpretative power in social science applications.

4 Summary of our research experience

4.1 The role of genre in Parliament

In a paper in the second part of this issue (Németh et al., 2025), we analyzed speeches in the Hungarian Parliament containing the term 'Kárpát-medence' (Carpathian Basin) using structural topic modeling. The term 'Kárpát-medence' is one of the most significant concepts in Hungarian geopolitical thinking, which also served irredentist goals in the period between the two world wars. The concept has been gradually revived in recent decades and is now part of everyday political discourse and national identity building. The stake of our research is that the concept of the Carpathian Basin is not a neutral concept, as it is charged with geopolitical intentions and is consciously used by actors in political discourse.

Parliamentary speeches have fixed types, and these types show important pragmatic differences. They can be classified into two main groups, according to Parliament's dual function (legislative and control). Keynote speeches, speeches, and pre-agenda speeches usually have a legislative function, i.e., a representative function that reflects the main messages of the parties. Speeches with a control function are usually immediate questions, immediate answers, and two-minute speeches. Another important genre distinction is whether the text is pre-written or spontaneous. Finally, debates are also audience-oriented, as they take place in front of a real audience and a virtual audience of voters.

We found that, compared to the corpus as a whole, the Carpathian Basin sub-corpus is more likely to contain pre-written speeches, which have a legislative function and represent the main messages of the parties and are pre-written. In other words, the use of the term 'Carpathian Basin' is important for speakers when they want to represent their point of view, suggesting the strong ideological function of the term. Similarly, we identified a characteristic difference in the distribution of genres within the topics.

In our mixed-methods interpretation of the topics, we found that the importance of the speech function was most pronounced in the topics dealing with administrative and regulatory issues affecting Hungarians living beyond the borders. This topic turned out to be quite interactive and included controversial issues. For example, 'you' was among the most relevant terms, suggesting that there were often arguments about the topic, and 'don't be angry' was also a relevant term, suggesting even heated arguments ('don't be angry, but…' is a phrase that is commonly used in Hungarian when one party in a dispute disagrees with the other). The qualitative analysis also showed that the topic is often included in heated debates.

That is, according to our observation, topics can be characterized not only on a descriptive or semantic level but also on a pragmatic level. There are topics of a more performative nature concerning the function of Parliament as a context, where speech is part of political action and can therefore be approached as a speech act.

4.2 Speech acts in Parliament

The importance of speech acts was first pointed out by Austin (1962), who stressed the importance of the non-referential use of language and the many functions that utterances perform as part of interpersonal communication. Austin showed that many utterances do not convey information but have an action value; he called these utterances 'performatives' as opposed to 'constatives.' From this perspective, the three components of a speaker's utterance are locution (the semantic meaning of the utterance), illocution (the speaker's intention, the performative function), and perlocution (how the listener received it). This approach has also fertilized the social sciences; see, for example, Habermas (1984).

Our experience shows that this approach can be fruitfully used to interpret the topic model. The interpretation can, in the first step, focus on the semantic level of the text: at the level of the topics, one can analyze substantive themes that refer to the locutionary level. At the same time, it is worth looking at word types separately in order to identify other types of speech acts. Adjectives refer partly to moods and partly to the illocutionary and perlocutionary levels. If we look at topics that include verbs as relevant terms, we can find descriptive, argumentative, or performative speech acts. A qualitative analysis of the most representative texts can support the classification.

As an illustration, we first cite our paper (Sik et al., 2024), where we set out to map the Hungarian parliamentary discourses of the last two decades from the perspective of democratization using structural topic modeling. During the interpretation, we identified both substantive characteristics (i.e., what topics and framing appear) and the pragmatic-level characteristics of the text (how the topic is discussed). After interpreting each topic individually using both quantitative and qualitative approaches, we classified the topics into three groups according to their performative nature. The three groups differed both in their typical word-class patterns and in their speech-act function. The first group (labeled 'descriptive speech acts') was constituted of policy-based topics discussing specific economic, social policy, and administrative issues, mostly in a technocratic, descriptive manner. The relevant terms associated with the topics in this group were nouns (e.g., expenditure, transfer, contractor, taxpayer, resource-sharing, micro-region, inflation, tax cut, livestock, agroeconomics), indicating a highly specialized expert language. However, these topics were not only similar to each other on a substantive level but also on a pragmatic level: the hermeneutic analysis of the exemplary speeches revealed that the tone of these speech acts was relatively rational compared to the other clusters. Debates revolving around expert issues tended to maintain a rational basis, as the parliamentarians did not evaluate each other but relied on discursive frames of justification.

The second group was labeled 'confrontational, argumentative speech acts.' It consisted of controversial topics such as crises, financial crimes, political scandals, or the neuralgic points of memory politics. The group was characterized by emotionally charged adjectives and attributes (e.g., heroic, noble, authoritarian, beautiful, illegal). As the hermeneutic analysis of the speeches of this group revealed, the tone of these speeches was highly emotional: the MPs did not rely on the logic of justification but instead gave emotionally heated, divisive speeches.

The third group contained most of the verbs within their relevant terms; this cluster consisted mostly of performative speech acts that served an illocutionary function. According to the hermeneutic analysis of the speeches, the common core of these substantively diverse speech act patterns was performative functionality: linguistic patterns expressing hostility (terms such as 'lying,' 'lies,' 'stupid,' 'ashamed') and politeness (e.g., 'respected,' 'greet,' 'ask,' 'honors,' 'kindly') defined this cluster. In contrast to the neutral speech acts of the first group and the evaluative speech acts of the second group, the third group was characterized by performative speech acts with an illocutionary function. This means that the topics of the third group showed how actors interact with each other while discussing substantive and divisive issues.

We then analyzed changes in the frequency of these three clusters over the last two decades, discovering that the frequency of speech acts expressing the formal rules of parliamentary practice has gradually declined along with the performatives expressing democratic civic culture, while indicators of abusive communication have increased dramatically. In sum, the pragmatic analysis that complements the descriptive semantic level added a lot to the message of our paper. Note that this pragmatic approach to interpreting the topic model is so unusual that one reviewer of our paper even commented that a dictionary-based approach or supervised learning might be more effective for investigating the confrontational or emotional nature of texts.

4.3 Speech acts in online forums about depression

We also cite here another study based on the study of online forums because, in this context, interaction, i.e., pragmatics, plays a prominent role – as also shown in the interpretation of the topics.

We have analyzed online forums about depression in several articles (Németh et al., 2020; 2021; 2022; Sik et al., 2023). First, we would like to briefly refer to a paper that used supervised learning, i.e., not a topic model (Németh et al., 2022), just to prove the interpretative importance of the word types mentioned above. Here, the goal of the machine learner was to classify the posts according to whether they frame the mental state using a biomedical, psychological, or social approach. In the interpretation of the classification, we saw that the role of the word types was important, e.g., texts with more nouns (i.e., more formal in tone according to linguistic research) were more likely to be defined as biomedical framing – this is presumably due to the more objective nature of the approach.

In our research that analyzed forums using the LDA topic model (Németh et al., 2021; Sik et al., 2023), the first dividing line in the interpretation of the topics was not drawn according to semantic differences but according to performative differences. Performative differences were the communicative function of the posts and the illocutionary and perlocutionary dimensions of the speech act. Posts were interpreted as monologues that addressed different aspects of the world and the self, and as interactions that focused on exchange with others. Within these categories, further subtypes were distinguished: monologues included more objective attributions and emotionally charged self-disclosures (both dominated by locutionary content); interactions included more pragmatic consultations (dominated by locutionary content) and quasi-therapeutic engagements (dominated by perlocutionary content). Differentiating between these patterns at the illocutionary level proved to be particularly important for mapping the complex discursive processes of online depression forums. In addition to discussing the challenges of a depressed life and sharing experiences and available countermeasures, these platforms also host 'ritual healing' attempts performed by engaged helpers who have already struggled with depression. In order to grasp these more complex potentials, the hermeneutic analysis of the topics proved to be essential.

Regarding the final model, the two main dividing lines between topics were drawn according to their substantive content and communicative functions. A methodologically noteworthy result was that this interpretation was also consistent with the distance map of topics generated by the LDAvis function (Mabey, 2020), where topics are plotted on a two-dimensional plane; this consistency showed the robustness of our results. In the later paper (Sik et al., 2023), we took an explicitly dual approach to the interpretation of topics, defining both their semantic content and their pragmatic functionality.

5 Conclusions

There are several different versions of the topic model, and the articles in this special issue include some of them. The conclusions of this article can be applied to all these variants.

What 'topics' are identified in the topics? The examples above show that in making sense of topics, it is worth focusing on two things at once: the semantic side of the texts (what are they talking about?) and the pragmatic side (how, for what purpose, with what effect on the audience?). We have also seen that it is worth applying speech act theory here, especially in social science research, where utterances are typically not just information carriers but actual actions that perform a variety of functions and have an impact on the speaker's environment as part of interpersonal communication. We have seen that in modeling the topics of online forums about depression, the topics can be interpreted along these two dimensions: the content of the posts and their communicative function. Similarly, in the study of Hungarian parliamentary speeches, the consideration of the genre and speech function of the texts has proved to be an important aspect.

Our case studies are based on less bounded genres (mostly spontaneous parliamentary speeches and online forums), where the importance of pragmatics is greater, but we believe that the majority of social research applications are based on just such informal genres. It is likely that the early and very widespread applications of the model did not emphasize this pragmatic aspect because they were mostly based on formal and written texts, such as newspaper articles (DiMaggio et al., 2013) or academic papers (Lafferty & Blei, 2009), so this aspect did not arise.

There are also methodological implications. On the one hand, we have drawn attention to the importance of paying attention to genre: genre-like information is available as metadata for many corpora.

On the other hand, we have seen how important the role of word forms is for speech acts – a result that highlights the risk of preprocessing decisions, such as word form filtering. In general, topic modeling is sensitive to the choice of preprocessing techniques (Grimmer & Stewart, 2013), and our analysis has shown why, for example, verb removal, which is a common procedure in general (Craswell et al., 2009) and in topic models in particular (Gautrais et al., 2017; Zirn & Stuckenschmidt, 2014), can lead to serious losses in terms of depth of interpretation and model validity.

References

Austin, J. L. (1962). *How to do things with words.* Oxford University Press.

Blei, D. M., Ng, A. Y. & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research, 3,* 993–1022.

Blevins, C. (2010, April 1). *Topic modeling Martha Ballard's diary.* https://www.cameronblevins.org/martha-ballards-diary/

Brookes, G. & McEnery, T. (2019). The utility of topic modelling for discourse studies: A critical evaluation. *Discourse Studies, 21*(1), 3–21. https://doi.org/10.1177/1461445618814032

Chakrabarti, P. & Frye, M. (2017). A mixed-methods framework for analyzing text data: Integrating computational techniques with qualitative methods in demography. *Demographic Research, 37,* 1351–1382. https://doi.org/10.4054/demres.2017.37.42

Craswell, N., Demartini, G., Gaugaz, J. & Iofciu, T. (2009). L3S at INEX 2008: Retrieving entities using structured information. In *Advances in Focused Retrieval: 7th International Workshop of the Initiative for the Evaluation of XML Retrieval, INEX 2008, Dagstuhl Castle, Germany, December 15–18, 2008. Revised and Selected Papers 7* (pp. 253–263). Springer.

Denny, M. J. & Spirling, A. (2018). Text preprocessing for unsupervised learning: Why it matters, when it misleads, and what to do about it. *Political Analysis, 26*(2), 168–189. https://doi.org/10.1017/pan.2017.44

DiMaggio, P., Nag, M. & Blei, D. (2013). Exploiting affinities between topic modeling and the sociological perspective on culture: Application to newspaper coverage of US government arts funding. *Poetics, 41*(6), 570–606.

Firth, J. (1957). A Synopsis of Linguistic Theory 1930–1955. In *Studies in Linguistic Analysis, Philological Society* (pp. 1–31). Oxford.

Gautrais, C., Cellier, P., Quiniou, R. & Termier, A. (2017). Topic signatures in political campaign speeches. In M. Palmer, R. Hwa & S. Riedel (Eds.), *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 2342–2347.)

Gillings, M. & Hardie, A. (2023). The interpretation of topic models for scholarly analysis: An evaluation and critique of current practice. *Digital Scholarship in the Humanities, 38*(2), 530–543. https://doi.org/10.1093/llc/fqac075

Grimmer, J. (2010). A Bayesian hierarchical topic model for political texts: Measuring expressed agendas in senate press releases. *Political Analysis, 18*(1), 1–35. https://doi.org/10.1093/pan/mpp034

Grimmer, J. & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis, 21*(3), 267–297. https://doi.org/10.1093/pan/mps028

Grondin, J. (2015). The Hermeneutical Circle. In N. Keane & C. Lawn (Eds.), *A Companion to Hermeneutics* (pp. 299–305). Wiley. https://doi.org/10.1002/9781118529812.ch34

Habermas, J. (1984). *The theory of communicative action*. Beacon Press.

Jacobi, C., Van Atteveldt, W. & Welbers, K. (2016). Quantitative analysis of large amounts of journalistic texts using topic modelling. *Digital Journalism, 4*(1), 89–106. https://doi.org/10.1080/21670811.2015.1093271

Jacobs, T. & Tschötschel, R. (2019). Topic models meet discourse analysis: a quantitative tool for a qualitative approach. *International Journal of Social Research Methodology, 22*(5), 469–485. https://doi.org/10.1080/13645579.2019.1576317

Lafferty, J. & Blei, D. (2009). Topic Models. In A. N. Srivastava & M. Sahami (Eds.), *Text Mining* (pp. 71–93). Chapman and Hall/CRC.

Laureate, C. D. P., Buntine, W. & Linger, H. (2023). A systematic review of the use of topic models for short text social media analysis. *Artificial Intelligence Review, 56,* 14223–14255. https://doi.org/10.1007/s10462-023-10471-x

Light, R. & Cunningham, J. (2016). Oracles of peace: Topic modeling, cultural opportunity, and the Nobel Peace Prize, 1902–2012. *Mobilization: An International Quarterly, 21*(1), 43–64. https://doi.org/10.17813/1086-671x-20-4-43

Mabey, B. (n. d.). *pyLDAvis: Python library for interactive topic model visualization. Port of the R LDAvis package.*

Mikolov, T., Sutskever I., Chen K., Corrado, G. S. & Dean J. (2013). Distributed Representations of Words and Phrases and their Compositionality. In *Proceedings of the Conference on Advances in Neural Information Processing Systems 26 (NIPS).* https://doi.org/10.48550/arXiv.1310.4546

Mützel, S. (2015). Structures of the Tasted: Restaurant Reviews in Berlin Between 1995 and 2012. In A. B. Antal, M. Hutter, & D. Stark (Eds.), *Moments of Valuation: Exploring Sites of Dissonance* (pp. 147–167). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198702504.003.0008

Németh, R., Katona, E., Balogh, P., Rakovics, Zs. & Unger, A. (2025). What else comes with a geographical concept beyond geography? The renaissance of the term 'Carpathian Basin' in the Hungarian Parliament. *Intersections. East European Journal of Society and Politics, 11*(1), forthcoming.

Németh, R., Máté, F., Katona, E., Rakovics, M. & Sik, D. (2022). Bio, psycho, or social: supervised machine learning to classify discursive framing of depression in online health communities. *Quality & Quantity, 56*(6), 3933–3955. https://doi.org/10.1007/s11135-021-01299-0

Németh, R., Sik, D. & Katona, E. (2021). The asymmetries of the biopsychosocial model of depression in lay discourses – Topic modelling online depression forums. *SSM-Population Health*, *14*, 100785. https://doi.org/10.1016/j.ssmph.2021.100785

Németh, R., Sik, D. & Máté, F. (2020). Machine learning of concepts hard even for humans: The case of online depression forums. *International Journal of Qualitative Methods, 19*, 160940692094933. https://doi.org/10.1177/1609406920949338

Röder, M., Both, A. & Hinneburg, A. (2015). Exploring the space of topic coherence measures. *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining – WSDM '15,* 399–408 https://doi.org/10.1145/2684822.268532

Schmiedel, T., Müller, O. & vom Brocke, J. (2019). Topic modeling as a strategy of inquiry in organizational research: A tutorial with an application example on organizational culture. *Organizational Research Methods, 22*(4), 941–968. https://doi.org/10.1177/1094428118773858

Shadrova, A. (2021). Topic models do not model topics: epistemological remarks and steps towards best practices. *Journal of Data Mining and Digital Humanities, 2021.* https://doi.org/10.46298/jdmdh.7595

Sik, D., Rakovics, Zs. & Németh, R. (2024). *Towards a culture of disrespect – topic modelling Hungarian parliamentary discourses.* Manuscript.

Sik, D., Németh, R. & Katona, E. (2023). Topic modelling online depression forums: beyond narratives of self-objectification and self-blaming. *Journal of Mental Health, 32*(2), 386–395. https://doi.org/10.1080/09638237.2021.1979493

Sievert, C. & Shirley, K. (2014). LDAvis: A method for visualizing and interpreting topics. *Proceedings of the Workshop on Interactive Language Learning, Visualization, and Interfaces* (pp. 63–70). Association for Computational Linguistics. https://doi.org/10.3115/v1/W14-3110

Weber, M. (1949). "Objectivity" in Social Science and Social Policy. In E. A. Shils & H. A. Finch (Eds.), *Max Weber on the methodology of the social sciences* (pp. 50–112). The Free Press.

Zirn, C. & Stuckenschmidt, H. (2014). Multidimensional topic analysis in political texts. *Data & Knowledge Engineering, 90*, 38–53. https://doi.org/10.1016/j.datak.2013.07.003