

Document-based Topic Coherence Measures for News Media Text

Damir Korenčić^{a,**}, Strahil Ristov^a, Jan Šnajder^{b,*}

^a*Department of Electronics, Rudjer Bošković Institute,
Bijenička cesta 54, 10000 Zagreb, Croatia*

^b*University of Zagreb, Faculty of Electrical Engineering and Computing,
Text Analysis and Knowledge Engineering Lab, Unska 3, 10000 Zagreb, Croatia*

Abstract

With the widespread and increased consumption of online news, there is a rising need for automated analysis of news text. Topic models have proven to be useful tools for unsupervised discovery of topics from large amounts of text, including news media texts. Topics produced by a topic model are often represented as probability-weighted word lists, and it is expected that these bear correspondence to *semantic topics* – semantic concepts representable by a topic model. However, because the quality of topics varies and not all topics correspond to semantic topics in practice, much research effort has been devoted to automated evaluation of topic models. One class of popular and effective methods focuses on *topic coherence* as a measure of a topic’s semantic interpretability and its correspondence to a semantic topic. Existing topic coherence methods calculate the coherence score based on the semantic similarity of topic-related words. However, news media texts revolve around specific news stories, giving rise to many contingent and transient topics for which topic-related words tend to be semantically unrelated. Consequently, the coherence of many news media topics is not amenable to detection via state-of-art word-based coherence measures. In this paper, we propose a novel class of topic coherence methods that estimate topic coherence based on topic documents rather than topic words. We evaluate the proposed methods on two datasets containing topics manually labeled for document-based coherence, derived from US and Croatian news text corpora. Our best-performing document-based coherence measure achieves an AUC score above 0.8, substantially outperforming a strong baseline method and state-of-art word-based coherence methods. We also demonstrate that there may be benefit in combining word- and document-based coherence measures. Lastly, we demonstrate the usefulness of document-based coherence measures for automated topic discovery from news media texts.

*Principal corresponding author

**Corresponding author

Email addresses: damir.korencic@irb.hr (Damir Korenčić), strahil.ristov@irb.hr (Strahil Ristov), jan.snajder@fer.hr (Jan Šnajder)

Keywords: Topic models, Topic coherence, Topic model evaluation, Text analysis, News text, Explorative analysis

1. Introduction

News media, including broadcast, the press, and online news, in many ways mold our perception of the world and influence our decisions – they are, as McCombs (2014) put it, paraphrasing Lippmann (1922), “our windows to the vast world beyond direct experience”, which “determine our cognitive maps of that world.” While the role of news media has not changed much since the days of Lippman, the technology certainly has. According to a recent study by Newman et al. (2016), online news, including online news sites, news aggregates, search engines, social media, and increasingly also messaging apps is now the predominant source of news. According to the same study, two-thirds of consumers prefer news in text, while a study by American Press Institute (2015) found that reading news from social media is woven into the daily lives of two-thirds of American millennials. Newman et al. (2016) also report that the majority of consumers – especially smartphone users and young adults – now discover news stories through algorithms rather than editors or journalists.

The widespread and increased consumption of online news in textual form has paralleled a growing interest in the use of natural language processing (NLP) and machine learning for automated analysis of news texts. Such technologies target diverse components of the news media ecosystem, from providing end-consumers with more efficient and personalized access to news (Steinberger et al., 2013; Vossen et al., 2014; Li et al., 2011) to support for news production and dissemination (Clerwall, 2014; Popescu & Strapparava, 2017) to content analysis in social science (Flaounas et al., 2013; Neuendorf, 2016). The typical NLP tasks include text categorization (Masand et al., 1992; Rose et al., 2002), exploratory analysis (Grobelnik & Mladenic, 2004; Chuang et al., 2014; Shahaf & Guestrin, 2012), news recommendation (Bogers & Van den Bosch, 2007; Garcin et al., 2013; Li et al., 2014), summarization (Gao et al., 2012; Lee et al., 2003), information extraction (Piskorski & Yangarber, 2013), topic detection and tracking (Allan, 2012; Phuvipadawat & Murata, 2010), event extraction (Tanev et al., 2008), and sentiment analysis (Godbole et al., 2007; Balahur et al., 2013; Nassirtoussi et al., 2015).

One important NLP task is the unsupervised discovery of topics from large volumes of texts, including news corpora. *Topic models* (Blei et al., 2003; Blei, 2012) have proven to be extremely useful tool for this task. A topic model is a probabilistic model of text that, given a set of text documents as input, produces word-topic and document-topic probability distributions – each topic is characterized by a probability-weighted word list while each document is characterized by probabilities of occurrence of topics in the text. Topics produced by a topic model are expected to correspond to concepts and can be used as topical summaries or semantic features for various downstream NLP tasks. Table 1 shows an example of topics produced by running a topic model on a US

news corpus. In addition to exploratory analysis (Newman et al., 2006; Shahaf & Guestrin, 2012; Chuang et al., 2014), topic models have also been applied to various other news-related tasks, such as news recommendation (Garcin et al., 2013; Li et al., 2014), summarization, (Gao et al., 2012), retrieval (Yi & Allan, 2009), and content analysis (Jacobi et al., 2016).

The main advantage of topic models is that they are unsupervised and require a minimum of linguistic processing (merely tokenization and optionally stopword removal and morphological normalization). However, the downside – pretty much as with all unsupervised approaches – is that the quality of the results may vary. What this means in practice is that not all topics produced by a topic model will correspond to *semantic topics* – the semantic concepts that are in principle representable by a topic model.¹ Some topics will clearly correspond to semantic topics, while others will be less recognizable, while yet others will simply be too noisy (Chuang et al., 2013; Roberts et al., 2016). The quality of the produced topics depends on a number of factors, including the properties of the corpus, the model structure, and its hyperparameters (Chuang et al., 2013). On top of that, the problem is compounded by the stochastic and approximate nature of the topic model inference.

The practical importance of topic models on one hand and the variability of model quality on the other have motivated extensive research on automated topic model evaluation. The approaches to topic model evaluation may be divided into *extrinsic* and *intrinsic*. The former approach is used to evaluate the quality of a model in terms of how much it improves the performance on a downstream NLP task, such as information retrieval (Wei & Croft, 2006), word sense disambiguation (Boyd-Graber et al., 2007a), sentiment analysis (Titov & McDonald, 2008), or word similarity and document classification (Stevens et al., 2012). In contrast, the intrinsic approach evaluates the quality of the produced topics irrespective of an application. Various metrics of model fit (Wallach et al., 2009), such as perplexity and held-out likelihood, have been proposed to evaluate the predictive accuracy of topic models. While certainly indicative of model quality, Chang et al. (2009) note that these metrics fail to address the exploratory aspect of topic models, which concerns the *semantic interpretability* of topics. More specifically, Chang et al. showed that models that fare better in predictive perplexity often have less interpretable topics, suggesting that evaluation should consider the internal representation of topic models and aim to quantify their interpretability. The idea soon gave rise to a new family of methods (Newman et al., 2010; Mimno et al., 2011; Lau et al., 2014; Röder et al., 2015; Ramrakhiani et al., 2017) that evaluate the semantic interpretability by measuring the *topic coherence*. These methods assume that topic coherence correlates with the coherence of the words assigned to that topic and thus quantify topic coherence as the coherence of top-ranked topic words.

¹We acknowledge here that not all semantic concepts can be modeled by topic models, i.e., represented as a list of words or documents. A case in point are the non-topical concepts, such as “good news”, “bad news”, “breaking news”, etc.

In other words, coherence of a set of words is taken as a proxy for semantic interpretability.

While the assumption that topic coherence correlates well with coherence of the top-ranked topic words is intuitive and certainly also true in many cases, we argue that it nonetheless provides a partial view of the notion of topic coherence. In particular, we note that word-topic distribution constitutes just a subset of the topic-related information contained in the model. Another important piece of information, thus far mostly overlooked in topic coherence analysis, is the document-topic distribution. It seems sensible that a coherent topic should be coherent in terms of not only words but also the documents it is associated with. Furthermore, it appears that quantification of topic coherence in terms of the associated documents rather than words should be more reliable, as documents contain more information than words.

Examples from Table 1 will serve to illustrate our point. The first two topics, labeled “Economy” and “Sport”, are quite interpretable based on the top-10 topic words; moreover, these words indeed appear coherent in terms of semantic relatedness. However, the two other topics, “US DHS shutdown” and “ISIS war authorization”, are quite different. Unlike “Economy” and “Sport”, which are general and enduring topics, these two topics are contingent and transient – a trait typical of topics from a newspaper corpus. The contingency of the topics manifests itself in topic words being less similar to each other in terms of semantic relatedness. Consequently, expert knowledge is required to estimate the coherence of these topics – in this case, the knowledge about the events and stories covered by the US news media in 2015. For example, the “US DHS shutdown” can be interpreted and labeled as such only if one is familiar with the political clash that occurred in February 2015 (involving the Speaker of the US House of Representatives John A. Boehner and Senator Mitchell McConnell Jr.), to the extent that one can discern this event from two other political events involving the same politicians later that year. Similarly, the “ISIS war authorization topic” could easily be misinterpreted as being about some other bill or resolution involving US Senate Chairman Bob Corker (for instance, the Iran Nuclear Agreement Review Act). These examples clearly illustrate that, for the transient topics corresponding to news stories the topic words provide insufficient information to assess the coherence and, equivalently, semantically interpret and label the topics. In such cases, however, the topic coherence can often be more easily established by inspecting the documents associated with the topic.

The focus of this paper is the automated topic model evaluation for news media texts. Motivated by the above observations, and in contrast with existing work, we approach the task by considering the coherence of topic *documents* rather than the coherence of topic words. We believe that document-based topic coherence can better capture topics’ semantic interpretability, especially when the topics are transient and contingent, as is often the case with news topics. The main result of our work is a novel method for calculating topic coherence. The method consists of three steps: (1) the selection of topic-related documents, (2) document vectorization, and (3) computation of a coherence

Topic label	Top-10 topic words
1. Economy	rate, economy, growth, fed, dropped, low, market, reserve, price, unemployment
2. Sport	team, game, players, season, sports, league, fans, football, bowl, pick
3. US DHS shutdown	boehner, homeland, block, dhs, mcconnell, pass, illegal, speaker, border, deportation
4. ISIS war authorization	ground, veto, resolution, corker, latino, bob, draft, review, capitol, pass

Table 1: Example topics derived from a corpus of news texts compiled by Korenčić et al. (2015), comprising about 24 thousand US news articles from 2015. Each topic is characterized by ten words with the highest word-topic probability. The topic labels were assigned manually based on inspection of documents with highest document-topic probability.

score from the document vectors using either distance-based, graph-based, or density-based methods. We consider a number of options for each of the three steps and experimentally evaluate the best variants on two datasets containing topics with manually annotated document-based coherence scores. We show that a graph-based coherence measure performs the best, outperforming a strong baseline document-based method.

Furthermore, we experiment with measuring document-based coherence using state-of-art word-based coherence measures, demonstrating that these measures fail to estimate document-based coherence. A qualitative analysis of topics based on word- and document-based coherence reveals that document- and word-based measures can complement each other and that it therefore may be beneficial to combine these two types of measures.

Lastly, in a proof-of-concept study, we demonstrate the usefulness of document-based coherence measures for the task of semi-automated topic discovery. In summary, the contribution of our work is threefold: (1) we introduce the notion of document-based topic coherence and demonstrate its adequacy for news media texts, (2) we propose a novel, document-based topic coherence measure, and (3) we compile and make available two datasets of topics manually annotated with document-based topic coherence, as well as the code and the resources necessary to replicate our experiments.²³

The remainder of the paper is set out as follows. The next section provides background on topic models and an overview of the related work, including applications on news texts and topic model evaluation. In Section 3, we elaborate the notion of the document-based coherence and propose methods for computing document-based coherence measures. In Section 4, we evaluate and analyze

²<https://rebrand.ly/doc-coh-dataset>

³<https://rebrand.ly/doc-coh-code>

the document-based coherence measures, while in Section 5, we compare them against state-of-art word-based coherence measures. In Section 6, we describe the proof-of-concept for the application of document-based coherence measures to semi-automated topic discovery. In Section 7, we conclude the paper and outline future work.

2. Background and Related Work

In this section, we give a brief description of topic models, followed by an overview of related work. There are three threads of research relevant to our work: applications of topic models for news text analysis, evaluation of topic models, and topic coherence evaluation.

2.1. Topic Models

Topic models (Blei et al., 2003) are generative probabilistic models of text with numerous text analysis applications, including exploratory analysis of text collections (Grimmer, 2009; Chuang et al., 2012), information retrieval (Wei & Croft, 2006), feature extraction (Chen et al., 2011), and natural language processing tasks, such as word sense disambiguation (Boyd-Graber et al., 2007b) and sentiment analysis (Lin & He, 2009). The structure of a topic model is defined by a set of random variables and relationships among them, which together define the probabilistic process of text generation. Typically, the variables of interest are topics, defined as probability distributions over words in the dictionary, and document-topic distributions, defining topic salience within each of the documents.

The most widely used topic model is Latent Dirichlet Allocation (LDA), proposed by Blei et al. (2003). This model posits a fixed number of topics, K , with each topic being a probability distribution over words in the dictionary. Topics are represented by a word-topic probability matrix ϕ , with ϕ_{ij} being the probability of word j in topic i . Similarly, documents are represented by a document-topic probability matrix θ , with θ_{ij} being the probability of topic j in the i -th document. The process of text generation unfolds as follows. First, each topic ϕ_i is sampled as a multinomial distribution from a Dirichlet prior distribution with parameter $\vec{\beta}$. Then, for each document D_i , θ_i is sampled as a multinomial distribution from a Dirichlet prior distribution with parameter $\vec{\alpha}$. Lastly, for each word within the i -th document, a topic z_{ij} is sampled from θ_i , and then the word is sampled from the topic $\phi_{z_{ij}}$. This generative process is summarized by the following probabilities for the word-topic matrix and the document collection:

$$p(\phi) = \prod_{i=1}^K \text{Dir}(\phi_i | \vec{\beta})$$

$$p(D_i) = \text{Dir}(\theta_i | \vec{\alpha}) \prod_j \text{Mult}(z_{ij} | \theta_i) \text{Mult}(w_{ij} | \phi_{z_{ij}})$$

Text documents $D = \{D_i\}$ are the observed variables of the model, and inference algorithms are used to estimate the word-topic distributions $p(\phi|D)$, document-topic distributions $p(\theta|D)$, and assignments of topics to words $p(z|D)$. Typically, the inference is performed using approximate inference methods, such as Gibbs sampling (Griffiths & Steyvers, 2004) and variational inference (Blei et al., 2003; Hoffman et al., 2010). Various extensions of the basic LDA model with a richer structure have been proposed in the literature, including those that model text document metadata (Mimno & McCallum, 2012) and relationships between topics (Blei & Lafferty, 2007), as well as models with a variable number of topics (Blei, 2012).

All the coherence measures considered in this article use either the word-topic probability matrix ϕ or the document-topic probability matrix θ to represent topics as lists of topic-related words or documents, respectively, and use this representation to compute the coherence of topics.

While LDA and its variants are certainly the most widely used topic models today, it should be noted that generative models are not the only approach to topic modeling. One alternative is matrix factorization models, such as latent semantic analysis (LSA) (Deerwester et al., 1990) and non-negative matrix factorization (NMF) (Lee & Seung, 1999). These models derive a set of latent factors by approximating the document-word matrix as a product of document-factor and factor-word matrices. The latent factors can be viewed as corresponding to topics, with the semantics of factors defined by document-factor or factor-word weights. Relevant for the work presented in this paper is the fact that coherence measures studied in this paper can also be applied to such factor-based topics, represented by either document-factor or factor-word weights.

2.2. Topic Models for News Text Analysis

Topic models have been applied to diverse news text processing tasks, ranging from exploratory analysis and scientific news analysis to commercially viable applications such as news recommendation, summarization, and retrieval.

Exploratory analysis. Chuang et al. (2014) demonstrate the potential of topic models for analysis of large-scale news collections by constructing topical signatures of news outlets and using them to characterize and compare the outlets. Newman et al. (2006) applied topic modeling and named entity recognition to a news corpus and demonstrated an analysis of topical trends and visualization of topic-entity and entity-entity relationships based on graphs constructed from topic model data. Ahmed et al. (2011) proposed a news-oriented topic model that models both topics and storylines and demonstrated the model’s application for exploratory analysis of news collections. Dou et al. (2012) proposed an interactive visualization system that combines event detection with topic modeling and named entity recognition as a means of providing rich contextual information for events.

Content analysis. Topic models have established themselves as a useful tool for quantitative content analysis within computational social science (Jacobi et al.,

2016), owing to the fact that they can scale up the analysis to large document collections and alleviate the labor-intensive process of manual document categorization (document coding). In the context of news media, two typical use cases are media agenda (McCombs & Shaw, 1972) and news framing (Entman, 1993) analyses. For instance, Kim et al. (2014) demonstrated the use of topic models for a comparative analysis of media agenda and public agenda issues. The analysis is carried out by first training a topic model on news articles and user comments and then comparing the salience of model topics corresponding to issues in news- and user-generated texts. Korenčić et al. (2015) addressed the problem of weak correspondence between model topics and news issues of interest and proposed a semi-supervised method for media agenda analysis consisting of news issue discovery and the measurement of issue salience. Jacobi et al. (2016) used topic models to first identify the topics corresponding to specific issues of interest and then analyzed how the framing of these issues has changed over time. Because topic models and other text mining tools are receiving increased interest in the social science community, Grimmer & Stewart (2013) pointed to a need for new methods for validating these tools before they can be adopted as standard. We see measures of topic coherence proposed in this paper as an important step toward that goal.

News recommendation. News recommendation algorithms address one of the key challenges of online news media: helping users find interesting news articles. Topic models have been used successfully for this purpose. Garcin et al. (2013) proposed a personalized news recommendation system based on a notion of context, defined as either a sequence of news items or a sequence or distribution of topics derived from a topic model. A probabilistic model of context change is defined, and recommendations are derived by combining several prediction models that use the context as input. Li et al. (2014) proposed a news recommendation system in which both user preferences and text clusters representing potential matches are constructed from topic distributions, and news items are recommended either by direct matching or by random-walk-based ranking.

Other applications. Topic models have been used for various other news text analysis tasks. For instance, Gao et al. (2012) proposed an event summarization method that uses a cross-collection topic model of news articles and tweets to produce summaries by ranking article sentences and tweets based on conditional probabilities derived from the model. Shahaf & Guestrin (2012) proposed a system for interactive discovery of news storylines, represented as temporal sequences of news articles connecting two endpoint articles, where each news article is represented using features constructed from topic models. Yi & Allan (2009) experimented with using topic models to enhance news information retrieval. Their method uses topic models for query expansion and the construction of a document language model, demonstrating performance increases in some use cases.

2.3. Evaluation of Topic Models

Topic models are only as useful as the quality of the topics they produce. As noted in the introduction, the downside of topic models is that the quality of the topics depends on a range of factors. First, topic modeling involves a number of design decisions, including which topic model structure and inference algorithm to use, how to set the model hyperparameters, and how to preprocess the text. Second, due to the stochastic nature of the inference process, the quality of the topics can greatly vary even for a single model. Methods for automated topic model evaluation help in addressing both issues: they can be used to narrow down the set of possible design options and to identify high-quality models from several runs.

In this paper, we focus on automated methods for intrinsic (i.e., task-independent) evaluation, which is more generally applicable than extrinsic (i.e., task-dependent) evaluation. The simplest and most common intrinsic evaluation method is to measure the perplexity of held-out text data with respect to the inferred model (Blei et al., 2003; Wallach et al., 2009). Perplexity is commonly used to demonstrate the viability of newly devised topic models.

Another measure of intrinsic model quality is *stability* (Waal & Barnard, 2008; Koltcov et al., 2014; Belford et al., 2018), defined as a measure of variance of the inferred models given a learning setup. Stability of a model is a desirable property in a number of applications, most notably in social sciences.

Mimno & Blei (2011) proposed a topic evaluation method based on measuring the discrepancy between empirical properties of the learned latent variables and properties expected from the probabilistic model structure. The method can be used to identify how well the topics fit the data.

Topic model evaluation can also be framed as a matching problem, in which model-produced topics are compared against ground-truth labels. In particular, (Ramirez et al., 2012) proposed the evaluation of topic models as clustering algorithms, by treating the topics as soft clusters and comparing them to a multi-labeled document collection using metrics for evaluating overlapping partitions. In contrast, (Chuang et al., 2013) proposed a framework for matching model topics directly to human-compiled concepts, identifying four types of misalignments between model-produced topics and reference concepts. In addition, they examined different topic models and model inference parameters with respect to the quality of topic-concept alignment.

Similarly to Ramirez et al. (2012), AlSumait et al. (2009) proposed a topic evaluation method that relies on documents from the corpus as external validation input. However, unlike Ramirez et al. (2012), they did not compare the topics against manually labeled documents, but rather score the topics based on topic-induced probabilities of words and documents. In particular, they computed the distances between topic-document or topic-word distributions on one side and uninformative distributions (uniform and “vacuous”) on the other and obtained the final score of topic quality by combining these two distances.

2.4. Topic Coherence

Relevant to the work described in this paper is the line of research that uses *topic coherence* as a measure of semantic interpretability and association with well-defined concepts (Newman et al., 2010). This line of research was motivated by the work of Chang et al. (2009), who framed the evaluation of topic coherence as a *word intrusion* task: human judges were asked to identify intruder words randomly inserted into a set of top topic words, with the idea that, for coherent topics, the intruders will be easier to spot. The intrusion-based scores were aggregated to obtain models’ scores, and the experiments showed that the human scores do not necessarily correlate with perplexity – a finding that indicated the need for new semantic measures of topic quality.

Subsequent approaches aimed to automate the computation of topic coherence scores. Most methods calculate the coherence of a topic by averaging pairwise semantic similarities of the most likely topic words (typically 5 or 10 words). The experimental evaluation is usually carried out using a ranking measure or a correlation coefficient to assess the agreement between predicted coherence scores and human-annotated coherence judgments (either binary coherent/non-coherent judgments or graded judgments on a Likert scale).

Newman et al. (2010) calculated coherence by averaging word similarity over pairs of top topic words, using WordNet-based and Wikipedia-based similarity metrics and pointwise mutual information of words. Another set of methods in (Newman et al., 2010) is based on the results of querying a web search engine with top topic words. Mimno et al. (2011) averaged the pairwise log-conditional probability of top topic words to calculate coherence and proposed a topic model learning algorithm that optimizes topic coherence. Musat et al. (2011) calculated topic coherence by mapping top topic words to WordNet concepts and finding optimal WordNet concepts that encompass the top words while being as specific (i.e., as low in the concept hierarchy) as possible. Aletras & Stevenson (2013) calculated coherence by averaging the similarity of distributional vectors, using normalized pointwise mutual information as the word-context weighting scheme, of top topic words over all the word pairs. Nikolenko et al. (2015) measured topic coherence by aggregating tf-idf based similarity of top topic words. O’Callaghan et al. (2015) used distributed word representations, also known as word embeddings (Mikolov et al., 2013), and calculated coherence by averaging cosine similarity of embedding vectors over pairs of top topic words. Nikolenko (2016) evaluated several variants of the coherence measure based on word embeddings, comparing them to existing coherence methods. The work of Lau et al. (2014) differs from the aforementioned work in that it first proposes a method to fully automate the original word intrusion task of Chang et al. (2009) and then compares the so-obtained coherence scores with similarity-based topic coherence scores.

Instead of calculating pairwise word similarities, Rosner et al. (2014) proposed the partitioning of the set of top topic words into subsets and averaged a “confirmation measure” based on conditional probability over the subset pairs. Ramrakhiani et al. (2017) proposed a coherence measure based on clustering

the embeddings of top topic words and approximating the coherence with the size of the largest cluster.

With the aim of facilitating a systematic evaluation of the many topic coherence measures proposed in the literature, Röder et al. (2015) proposed a generic framework for topic coherence computation. A generic topic coherence measure is defined in terms of how it partitions the set of top topic words into subsets, how the similarity of word sets is measured, and which method is used to derive word co-occurrence counts used for calculating the similarity. By searching through the space of measures defined in this manner, the authors derived several novel and efficient topic coherence measures.

Stevens et al. (2012) and O’Callaghan et al. (2015) reported on the use of topic coherence for empirical comparative analysis of different topic models. In particular, Stevens et al. (2012) compared the coherence of topics produced by LDA, NMF (Lee & Seung, 1999), and LSA (Deerwester et al., 1990) models and concluded that LDA and NMF topics are of comparable coherence and higher than those of LSA. Furthermore, NMF topics are shown to exhibit higher variance in coherence. O’Callaghan et al. (2015) investigated how pre- and post-processing may affect topic quality. Experiments showed that NMF models using tf-idf term weighting have higher mean coherences than standard LDA models. However, when term re-weighting is applied to LDA topics, their coherences became comparable to those of NMF. Experiments also showed that LDA tends to induce general topics, while NMF tends to induce more coherent and more specific topics.

The work presented in this paper falls under the category of topic coherence methods. However, our work differs from all of the above in that we measure topic coherence by reference to documents rather than words associated with a topic. In other words, the coherence measures we propose make use of document-topic distributions rather than word-topic distributions. As argued in the introduction, we hypothesize that characterizing topic coherence in terms of documents associated with the topics may be more adequate in some cases, especially for news media, which abounds with contingent and transient topics. Therefore, unlike all the prior work on topic coherence, we evaluate the proposed document-based coherence measures on datasets annotated specifically for document-based coherence, with topic coherence scores obtained upon inspection of topic-related documents rather than topic-related words. Technically, the work most similar to ours is that of AISumait et al. (2009), who proposed a measure of topic quality based on document-topic distribution and combined it with a word-based measure to produce the final topic quality score. However, unlike AISumait et al. (2009), we perform a quantitative evaluation of document-based measures using the document-based measure from AISumait et al. (2009) as a baseline. Our evaluation procedure bears similarities to that of Ramirez et al. (2012), as both evaluation approaches use document-topic distributions, but whereas Ramirez et al. (2012) evaluated complete topic models using manually labeled documents, we evaluate the individual topics by computing their coherence scores.

3. Document-based Topic Coherence Measures

In this section, we describe the proposed document-based topic coherence measures. The calculation of each measure comprises three main steps: (1) selection of topic-related documents, (2) vectorization of the documents, and (3) calculation of a coherence score based on document vectors.

The first step takes as input a topic and document-topic probability matrix θ and outputs a list of documents. Top documents are selected by taking a fixed number of documents with the highest document-topic probabilities. The document vectorization step takes as input a list of text documents and outputs a list of vectors. Vectorization of documents’ text is performed using standard bag-of-words or tf-idf vectorization or by aggregating the embeddings of document words. Finally, the vectorized documents are given as input to one of three types of methods for coherence scoring: distance-based, density-based, or graph-based methods. Distance-based methods aggregate pairwise document distances, while density-based methods use a vector probability density to approximate mutual closeness of documents. Graph-based methods derive a graph from document distances and calculate the coherence using one of several graph property measures. We next describe in detail each of the three steps of coherence measure calculation.

3.1. Selection of Topic-Related Documents

The aim of this step is to construct a list of text documents that are representative of a model topic – the documents that are associated with the topic in the context of the topic model. Selecting too many documents (in the extreme case, all the documents) will render the document list incoherent. Alternately, selecting too few documents (in the extreme case, a single top document for the topic) will likely make the list highly coherent.

We opt for a simple and model-independent strategy: given a topic, we select *TopDocs* documents with the highest document-topic probabilities. *TopDocs* is the parameter of the selection step and we optimize its value empirically. A similar strategy has proven effective in the case of word-based coherence measures, where selecting the top 10 words has shown to yield good results.

Recall from Section 2.1 that the document-topic distributions are represented by the document-topic probability matrix θ , where θ_{ij} is the probability of topic j in the i -th document. These probabilities represent the strength of association between the topics and the documents. Formally, for a topic j , the first *TopDocs* documents are chosen from the list of all documents D_{i_1}, \dots, D_{i_N} ordered by the probability of topic j in descending order ($\theta_{i_1 j} \geq \theta_{i_2 j} \geq \dots \geq \theta_{i_N j}$).

3.2. Document Vectorization

The purpose of the vectorization step is to transform information contained in documents’ texts into vectors that will be input to coherence scoring methods. We experiment with two standard text vectorization methods based on word occurrence counts: word probability and tf-idf. In addition, we experiment with per-document aggregation of two types of word embeddings derived from

a sizable external corpus: CBoW (Mikolov et al., 2013) and GloVe (Pennington et al., 2014).

Word count vectorization. These vectorization methods rely on counting occurrences of words in text documents. This is preceded by document preprocessing, which at the minimum includes tokenization and case-folding, but may also include morphological normalization, such as stemming or lemmatization.

Let N denote the total number of documents in the corpus, c_{ij} the number of occurrences of word j in the i -th document, d_i the size of the i -th document, and dc_j the number of documents the word j occurs in. Probability vector $prob_i$ of the i -th document is a vector of empirical word-in-document probabilities, obtained as maximum likelihood estimates, $prob_{i,j} = c_{ij}/d_i$. Tf-idf representation (Salton & Buckley, 1988) combines word-in-document probabilities with frequencies of word occurrence in other corpus documents. We use the tf-idf variant in which the tf-idf vector $tfidf_i$ of the i -th document is defined as $tfidf_{i,j} = tf_{i,j} \times idf_j$ where $tf_{i,j} = \log(c_{ij}) + 1$ and $idf_j = \log((N + 1)/(dc_j + 1)) + 1$. In addition, we normalize document tf-idf vectors to a unit L2-norm.

Word embedding aggregation. This method of aggregation relies on pre-constructed word embeddings (Turian et al., 2010): low-dimensional, continuous-valued vectorial representations of words’ meanings derived from word co-occurrences in a large text corpus. We experiment with the two most commonly used word embeddings: CBoW (Mikolov et al., 2013) and GloVe (Pennington et al., 2014). CBoW embedding vectors are obtained by optimizing the log-linear prediction of words based on their context words. On the English dataset, we use pre-trained 300-dimensional CBoW embeddings, derived from a 100-billion-word Google News corpus. On the Croatian dataset, we use the word2vec tool to train 300-dimensional CBoW embeddings on the large hrWaC corpus – a web corpus of Croatian texts (Ljubešić & Erjavec, 2011).⁴ GloVe embedding vectors are obtained by approximating global word co-occurrence probabilities using a weighted least squares regression model. On the English dataset, we use pre-trained 300-dimensional GloVe embeddings, derived from Wikipedia and Gigaword corpora. On the Croatian dataset, we train 300-dimensional GloVe embeddings on the hrWaC corpus using the glove tool.⁵

In accordance with standard practice (Wieting et al., 2015), we make use of the additive compositionality property of the word embeddings and compute the vector representation of a document text by adding up the word embedding vectors of all its content words. Optionally, we average the resulting vector to account for the differences in document lengths.

⁴CBoW pre-trained vectors for English and the word2vec tool are available from <https://code.google.com/archive/p/word2vec/>

⁵GloVe pre-trained vectors for English and the tool are available from <https://nlp.stanford.edu/projects/glove/>

3.3. Coherence Scoring Method

After documents representative of a topic have been selected and vectorized, a list of document vectors is fed to a coherence scoring method. We experiment with three types of methods: (1) distance-based methods, which aggregate distances between document vectors, (2) density-based methods, which approximate coherence using a multivariate normal distribution as a model of document vectors, and (3) graph-based methods, which first construct a connectivity graph from document vectors and then compute a suitable graph measure. Intuitively, these methods calculate coherence scores that correspond to mutual closeness, compactness, or connectivity among a set of document vectors.

3.3.1. Distance-based coherence

Distance-based methods rely on a measure of distance between vectors, i.e., a function $d : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^+$ that assigns a positive number to a pair of vectors. The distance measure does not have to be a metric in the mathematical sense, as long as it gives a useful notion of distance, such as the cosine distance.

We consider two simple distance-based methods: (1) average distance, which calculates an average of distances between all pairs of document vectors, and (2) distance variance, which calculates the average of distances between the document vectors and the center (mean) vector. In both cases, the final coherence score is produced by negating the average to convert a measure of dispersion into a measure of coherence.

3.3.2. Density-based coherence

The density-based coherence method works by first fitting a multivariate normal probability density function to the set of document vectors and then approximating coherence as the average log-density of the vectors under the model. The intuition is that, the higher the density, the tighter the grouping around the mode of the probability density function, and the higher the coherence of the document vectors.

The parameters of the multivariate density function are a mean vector, $\mu \in \mathbb{R}^n$, and a covariance matrix, $\Sigma \in \mathbb{R}^{n \times n}$. To reduce the number of parameters, and in turn prevent overfitting, we restrict the covariance matrix to either a diagonal matrix ($\Sigma = \text{diag}(\sigma_i^2)$, where σ_i^2 is the variance of the i -th vector component) or an isotropic matrix ($\Sigma = \sigma^2 \mathbf{I}$). This is equivalent to assuming uncorrelated noise and isotropic noise, respectively.

We fit the probability density function to the document vectors using maximum likelihood estimate. Before fitting, we optionally perform dimensionality reduction on the document vectors using the standard principal component analysis (PCA).

3.3.3. Graph-based coherence

Graph-based methods first construct a graph of selected topic-related documents and then calculate a coherence score using a suitable graph measure. The nodes of the graph correspond to the documents, while graph edges are

constructed using a measure of distance between document vectors. Graph measures we experiment with correspond to the notion of graph compactness or connectivity: closeness centrality, communicability centrality, clustering coefficient, number of connected components, and the size of the minimum spanning tree.⁶

Graph construction. We experiment with two methods to construct edges between document nodes. The first method constructs a fully connected weighted graph with edge weights set to the distance between document vectors. The second method uses a distance threshold to connect only those pairs of documents whose distance falls below a given threshold. After thresholding, the remaining edges can retain their weights, or the graph can be converted into an unweighted graph.

Closeness centrality. The first graph measure we consider is closeness centrality (Freeman, 1978). Closeness centrality of a graph node v is defined as the inverse of the average shortest path distance between the node and all other reachable nodes:

$$cc(v) = \frac{|C(v)| - 1}{\sum_{w \in C(v)} d(v, w)} \quad (1)$$

where $C(v)$ is the set of all nodes reachable from the node v (the nodes in the connected component containing the node v). Closeness centrality of an isolated node ($C(v) = v$) is 0.

To avoid assigning high closeness centrality to nodes of a fragmented graph (a graph with many small connected components), the closeness centrality is normalized by the relative size of the node’s connected component

$$cc_{\text{norm}}(v) = \frac{|C(v)| - 1}{N - 1} \frac{|C(v)| - 1}{\sum_{w \in C(v)} d(v, w)} \quad (2)$$

where N is the number of nodes in the graph.

We calculate the coherence score as the average normalized closeness centrality of all the graph nodes:

$$CC(G) = \frac{1}{N} \sum_{v \in G} cc_{\text{norm}}(v) \quad (3)$$

Subgraph centrality. Subgraph centrality (Estrada & Rodriguez-Velazquez, 2005) is a measure of node centrality correlated with the number of closed walks (cycles allowing node repetition) that start and end in a node. Let $\mu_k(v)$ denote

⁶All of the measures considered here are available as part of the NetworkX (Schult & Swart, 2008) library available at <http://networkx.readthedocs.io>.

the number of closed walks of length k originating in the node v . Subgraph centrality of a node v is defined as:

$$sc(v) = \sum_{k=1}^{\infty} \frac{\mu_k(v)}{k!} \quad (4)$$

The number of closed walks $\mu_k(v)$ is scaled by $k!$ to ensure the convergence of the series. Subgraph centrality of a node can be efficiently computed via spectral decomposition of the graph's adjacency matrix. Edge weights are irrelevant for subgraph centrality because the measure is based on the number of walks, not the weights of the walks. For this reason, subgraph centrality is not applied to complete graphs, as all the graphs with the same number of nodes would be assigned the same centrality score.

We calculate the coherence score by averaging subgraph centralities of all the nodes:

$$SC(G) = \frac{1}{N} \sum_{v \in G} sc(v) \quad (5)$$

Clustering coefficient. The clustering coefficient of a node v is the number of actual triangles that go through the node v , denoted as $T(v)$, divided by the number of all possible triangles that could go through that node. A triangle through a node v corresponds to a set of three distinct nodes – v , u_1 , and u_2 – such that edges vu_1 , u_1u_2 , and u_2v exist. The clustering coefficient is defined as:

$$cc(v) = \frac{T(v)}{\frac{deg(v)(deg(v)-1)}{2}} = \frac{2T(v)}{deg(v)(deg(v)-1)} \quad (6)$$

A weighted version of the clustering coefficient, which we apply to weighted graphs, is defined as (Saramäki et al., 2007):

$$cc(v) = \frac{1}{deg(v)(deg(v)-1)} \sum_{u_1, u_2} (w'(v, u_1)w'(u_1, u_2)w'(u_2, v))^{1/3} \quad (7)$$

Here, the sum is over all pairs of nodes that close a triangle with v , and $w'(u, v)$ is the weight of the edge between nodes u and v divided by the maximum edge weight in the graph.

We calculate the coherence score by averaging the clustering coefficients of all the graph nodes:

$$CC(G) = \frac{1}{N} \sum_{v \in G} cc(v) \quad (8)$$

Connected components and spanning tree. We experiment with two additional measures based on the connectivity structure of the graph. The first measure is the inverse of the number of connected components in the graph. We use

this measure only in combination with thresholded graph construction, as for complete graphs, the number of connected components is always one. The second measure is the negative weight of the minimum spanning tree of the graph. We use this measure only in combination with a non-thresholded, fully connected weighted graph.

4. Experiments with Document-Based Coherence

In this section, we evaluate and analyze the document-based coherence measures proposed in Section 3. To this end, we use two datasets of topics manually annotated for document-based coherence. After defining a set of coherence measures that correspond to sensible parameter values, we select the best measures on a development set and evaluate these on two test sets, using area under the ROC curve (AUC) as the performance metric. Lastly, we analyze the best-performing measures. We begin by describing the construction of the datasets.

4.1. Datasets

As argued in Section 1, document-based coherence has the potential to better capture the semantic interpretability of news topics than word-based coherence. Namely, as demonstrated in Table 1, assessing the coherence of some model topics requires the inspection of topic-related documents, as topic words themselves do not provide sufficient information. These observations motivate the use of a news domain dataset annotated specifically for document-based coherence.

To the best of our knowledge, no dataset that meets the above desiderata is publicly available; however, the dataset of Korenčić et al. (2015) is a good starting point. This dataset, built as part of a media agenda measuring study, contains LDA topics derived from a corpus of approximately 24k political news articles collected from news feeds of mainstream US news outlets between January 26th and April 13th, 2015. To increase the coverage of the semantic topics by model topics, Korenčić et al. (2015) ran five differently parametrized LDA models – three models with 50 topics and two models with 100 topics – and pooled the so-obtained topics, yielding a total of 350 topics. The topics were then manually labeled by two annotators (Korenčić et al., 2015). The annotators used an interface that displayed the top-ranked topic-related news articles, as well as the top-ranked topic words, and assigned zero or more labels to each topic. Each label corresponded either to one of the semantic topics or to a “noise” label indicating the presence of noise.⁷ The top-ranked documents were presented to the annotators as a list of titles sorted by document-topic probabilities, with the full texts of articles available for inspection. Similarly, topic-related words were presented as a list of words sorted by word-topic probabilities. For labeling, the annotators relied primarily on topic-related documents, as topic-related words proved either seemingly unrelated, vague, or too

⁷In Korenčić et al. (2015), the process was described as labeling of topics with *themes*, which can be taken to correspond to semantic topics.

abstract (cf. Tables 7 and 8 for example topics from our dataset), whereas lists of well-formed document titles provided more accurate and specific information. Consequently, the decisions on the semantic topics and their correspondence to model topics, as well as the decision on existence of noise, were made based on topic-related documents, while the words at best served to confirm this decision. In summary, of the 350 labeled model topics, 183 topics (52%) were labeled with one semantic topic, 52 topics (15%) were labeled with one semantic topic and the noise label, 61 topics (17%) were labeled with two semantic topics, 13 topics (4%) were labeled with two semantic topics plus noise, while 41 topics (12%) were labeled noise.⁸

For our evaluation, we adopt the dataset of Korenčić et al. (2015), with one straightforward modification. Because we are ultimately interested in coherence judgments for each of the topics rather than the topic labels, we use the human-provided topic labels to automatically generate binary coherence judgments. More specifically, we define as *coherent* the topics that have been annotated with a single semantic topic, possibly with the addition of noise, and as *incoherent* otherwise.⁹ In other words, we consider a topic as coherent as long as a human annotator can recognize that the topic corresponds to a single semantic topic, which is in line with the definitions of topic coherence used in Newman et al. (2010) and Mimno et al. (2011), whereas an incoherent topic corresponds either to noise or to a mixture of two or more semantic topics. After this procedure, 235 topics (67%) were labeled coherent, whereas 115 topics (33%) were labeled incoherent.

Tables 7 and 8 show a selection of model topics with the corresponding manually annotated semantic topics, obtained by the inspection of top-ranked topic-related documents. The tables also show the lists of top-ranked topic words for each topic. Similarly to the examples from Table 1, these examples illustrate that the set of top-ranked words is in most cases ambiguous and not informative enough to indicate the semantic topic.

We randomly split the described set of 350 topics into two subsets: the *development set*, containing 120 topics, and the *test set*, containing 230 topics. We use the development set to optimize the parameters of the coherence measures and the test set for final evaluation of these measures. To ensure that both the development and the test sets are representative of the entire dataset, the split was stratified across the five labels: single semantic topic, semantic topic plus noise, two semantic topics, two semantic topics plus noise, and noise.

⁸The proportions differ slightly from the ones in (Korenčić et al., 2015) because of an additional improvement step in which the annotators re-examined the model topics and corrected the labeling errors.

⁹Following (Nikolenko et al., 2015; Nikolenko, 2016), we work with binary judgments of coherence, since obtaining these from existing labels is rather straightforward. The alternative would have been to convert the existing labels into graded judgments of topic coherence, but it is not clear how one would proceed about this. Moreover, the use of binary labels paired with the AUC evaluation metric lends itself to a straightforward statistical interpretation, as detailed in Section 4.2.

In addition to the US news topics dataset described above, we introduce a second dataset that serves as an additional test set for assessing the robustness of the results. We refer to this second dataset as *test-cro*. The test-cro dataset consists of topics derived from a corpus of news texts in Croatian language, originally compiled for a media agenda setting study in (Korenčić et al., 2016). The topics were pooled from four LDA models – three models with 50 topics and one model with 100 topics, which resulted in a total of 250 topics.¹⁰ Coherence labels were derived by the same procedure used for the US topics – model topics were first annotated with semantic topics and the “noise” label, after which the coherence labels were derived from the annotations. Of the 250 topics, 166 topics (66%) were labeled coherent, whereas 84 topics (34%) were labeled incoherent.

We make available both the US news topics dataset and the Croatian news topics dataset.¹¹

4.2. Evaluation of Coherence Measures

We use the Area Under the ROC Curve (AUC) metric (Ling et al., 2003) to evaluate the performance of topic coherence methods. AUC is a metric employed for both classification and ranking evaluation. As a ranking measure, it has been used to evaluate word-based topic coherence (Nikolenko et al., 2015; Nikolenko, 2016) by comparing binary topic labels (coherent or incoherent) against the numerical coherence scores produced by the coherence measures.

Generally, given a model M producing confidence scores and data points $x \in D$ labeled with binary class labels, AUC corresponds to the probability that, for two points x and x' such that x belongs to the positive and x' belongs to the negative class, the positive one receives a higher score, i.e., $M(x) > M(x')$ (Nikolenko et al., 2015). In the case of coherence measures, with topics labeled as either coherent (positive class) or incoherent (negative class), AUC of a coherence measure Coh is the probability that, for a coherent topic t and an incoherent topic t' , the coherent topic gets a higher coherence score, i.e., $Coh(t) > Coh(t')$. The AUC scores are confined to the $[0, 1]$ interval, with 0 and 1 being the worst and the best score, respectively, and 0.5 being the expected score of an uninformative random measure.

The alternative interpretation of AUC rests on the idea that a model M producing confidence scores can be converted into a binary classifier by thresholding its output. The false positive rate (fall-out) and the true positive rate (recall) of the classifier for different threshold values define the receiver operating characteristics (ROC) curve. The performance of a perfect classifier corresponds to point $(0, 1)$ (no fall-out, complete recall). In contrast, the performance of a random classifier for different threshold values corresponds to a straight line from $(0, 0)$ to $(1, 1)$. Given an ROC curve, AUC is defined as the area under

¹⁰ Korenčić et al. (2016) originally used 200 topics, however, to make the size of the test-cro dataset comparable to that of the test set, we built and labeled an additional model with 50 topics.

¹¹<https://rebrand.ly/doc-coh-dataset>

the ROC curve. For the case of a binary classifier for topic coherence based on thresholding a coherence measure, the recall corresponds to the proportion of coherent topics detected by the classifier, while fallout corresponds to the proportion of incoherent topics falsely detected as coherent.

4.3. Baseline Method

We use as the baseline the document-based measure of “topic significance,” proposed by AlSumait et al. (2009). To the best of our knowledge, this measure is the only document-based measure of topic quality. The measure represents each topic as a probability distribution over the set of corpus documents obtained by normalizing topic-document probabilities. Topic significance is then calculated as the distance between the described distribution and the uninformative uniform distribution, using either cosine or KL-divergence as the distance measure. We use the variant based on the cosine distance, since it performs better on both datasets. In (AlSumait et al., 2009) this measure is not evaluated on its own but instead combined with similarly defined word-based measures in a composite measure of topic quality, which is then evaluated qualitatively by an inspection of high- and low-scored topics.

4.4. Model Selection

Our goal is to identify the well-performing document-based coherence measures from Section 3: those that have a high correlation with human-provided scores of document-based coherence of model topics, as measured by the AUC score. To this end, we first introduce a set of parameters that describe the structure of these measures. We then proceed to define a set of sensible parameter values corresponding to a set of coherence measures that will be considered in further evaluation.

Coherence measure categories. To ease the analysis, we group the coherence measures into six categories, as shown in Table 2. Each category corresponds to a pairing of two attributes: the coherence scoring method and the document vectorization method. We consider these two attributes to be the most distinguishing properties of a coherence measure.

The first attribute is the coherence scoring method, which essentially determines how the documents are viewed (as points in a vector space or as nodes in a graph) and how to estimate the coherence score from a set of documents (cf. Section 3.3). Distance-based methods (DISTANCE) rely on a measure of vector distance and aggregate the distances directly, while density-based methods (DENSITY) rely on a probability density defined on the space of vectors and calculate the dissipation of vectors around the center of distribution. In contrast, graph-based methods (GRAPH) add structure to the set of vectorized documents by constructing a document graph using a measure of vector distance for edge definition and calculate the coherence using measures describing various graph properties.

The second attribute, the document vectorization method (cf. Section 3.2), determines the representation of the top-ranked topic documents that are given

Scoring	Vectorization	# Measures
DISTANCE	CNT	48
DISTANCE	EMBD	80
DENSITY	CNT	96
DENSITY	EMBD	128
GRAPH	CNT	936
GRAPH	EMBD	1560

Table 2: Six categories of coherence measures, each corresponding to one combination of the coherence scoring and document vectorization methods, along with the number of distinct measures considered in the evaluation.

as input to the coherence scoring method. Here, we distinguish between two types of vectorization methods: CNT and EMBD. The former methods are based on word-in-document counts (normalized bag-of-words and tf-idf) derived from the same corpus that was used to build the topic model under evaluation. In contrast, EMBD vectorization methods refer to representations based on the aggregation of word embeddings (CBOW and GloVe), which have been derived from a large, external corpus. Apart from the difference in how the vectors are constructed in these two cases, an important difference is that the CNT vectors are domain-specific (in our case: the domains of US and Croatian political news), whereas EMBD vectors are generic (i.e., mixed-domain). This difference is likely to have an influence on the coherence measure calculation: in contrast to domain-specific vectors, the generic vectors will generally be more ambiguous and might not correspond to the domain-specific senses of some words.¹² Alternately, generic vectors might better capture the meaning of rare words, and might generally be statistically more reliable because they are derived from a larger corpus.

Coherence measure parameters. To allow for a systematic analysis of the document-based coherence measures proposed in Section 3, we parametrize these measures with respect to the selection of topic-representative documents, the vectorization method, and the coherence scoring method. Table 3 outlines the parameters and their values considered in subsequent experiments. The parameters are broken down by coherence scoring method except for the first three parameters, which are shared by all methods. The shared parameters together define how a model’s topic is transformed into a set of document vectors, while the method-specific parameters define the details of the coherence score calculation. For a more precise description of the parameters, the reader is referred to Section 3.

Note that not all parameter-value combinations are sensible. More specifi-

¹²As many polysemous words have domain-specific senses, restricting the domain from which the representations are derived will typically decrease the word-level ambiguity. This sense-domain relation has also been leveraged for improving word sense disambiguation, e.g., (Magnini et al., 2002).

Scoring	Parameter	Values
(All)	<i>TopDocs</i>	10, 25, 50, 100
	<i>DocVect</i>	bow, tfidf, cbow, glove
	<i>EmbeddingAgg</i>	average, sum
DISTANCE	<i>DistanceMeasure</i>	11, 12, cosine
	<i>DistanceAgg</i>	average, variance
DENSITY	<i>CovMatrix</i>	isotropic, diagonal
	<i>DimReduce</i>	None, 5, 10, 20, 50, 100
GRAPH	<i>DistanceMeasure</i>	11, 12, cosine closeness-centrality, subgraph-centrality,
	<i>GraphAlgo</i>	clustering, num-connected, min-spanning-tree
	<i>DistanceThresh</i>	None, 0.02, 0.05, 0.1, 0.25, 0.5, 0.75
	<i>Weighted</i>	True, False

Table 3: Coherence measure parameters and values for the different coherence scoring methods. The first three parameters are shared by all three scoring methods.

cally, *EmbeddingAgg* is applicable only if *DocVect* is **cbow** or **glove**. For DENSITY scoring method, if *DocVect* equals **bow** or **tfidf** (high-dimensional vectors), the value of *DimReduce* is varied among the full set of values (5, 10, 20, 50, 100), while if *DocVect* equals **cbow** or **glove** (vectors of lower dimension), the value of *DimReduce* is varied among the values 5, 10, and 20. The *DimReduce* value of **None** is combined with all the vectorization methods, resulting in a vector size of approximately 24k dimensions for **bow** and **tfidf**, and 300 dimensions (the original size of the word embeddings) for **cbow** and **glove**. Similarly, not all parameter combinations make sense for the GRAPH scoring method: for fully connected weighted graphs (*DistanceThresh* set to **None**), only the **closeness-centrality**, **clustering**, and **min-spanning-tree** methods are sensible. For thresholded graphs (*DistanceThresh* set to a positive real number), **subgraph-centrality** and **num-connected** are used only for unweighted graphs (*Weighted* set to **False**). Table 2 lists the numbers of sensible parameter combinations for each of the measure categories. In total, we consider 2,848 distinct coherence measures.

Another point worth mentioning is the treatment of the *DistanceThresh* parameter for the GRAPH scoring method. To account for the fact that different distance measures and vectorization methods generally yield distances at different scales, we proceed as follows. For each combination of *DistanceMeasure* and *DocVect*, we estimate the distribution of distances from a random sample of 100k document-vector pairs from the corpus. We then treat *DistanceThresh* as a percentile rank from the so-obtained distribution, varying the threshold among the values 0.02, 0.05, 0.1, 0.25, 0.5, and 0.75.

Measure Category		test		test-cro	
Scoring	Vectorization	AUC	p-value	AUC	p-value
GRAPH	CNT	0.804	–	0.812	–
DISTANCE	CNT	0.754	0.001	0.785	0.028
DENSITY	CNT	0.745	0.000	0.774	0.009
DISTANCE	EMBD	0.732	0.001	0.746	0.029
<i>doc-dist-cosine</i>	–	0.730	0.006	0.748	0.025
DENSITY	EMBD	0.728	0.001	0.725	0.005
GRAPH	EMBD	0.694	0.003	0.671	0.000

Table 4: Coherence measures’ AUC scores for each of the six categories and the baseline, ordered by the score on the test set. The p-values are derived by comparing the AUC score from the first row with the AUC score from the other six rows.

4.5. Results

Table 4 shows the AUC scores on the test and test-cro sets for each category of coherence measures. For each category, we report the AUC score of the best-performing coherence measure from that category on the development set. The best-performing measures are chosen from the set of measures defined by parameters described in Section 4.4. Note that we carry out no such optimization on the test-cro dataset, as we want to test the cross-dataset robustness of the measures’ parameters. The *doc-dist-cosine* on both datasets is the baseline method (cf. Section 4.3).

As seen from Table 4, the best-performing measure comes from the GRAPH-CNT category, which achieves an AUC score of over 0.8 on both test and test-cro sets and outperforms other measures by at least 0.027 AUC. We use the DeLong’s test (DeLong et al., 1988)¹³ to test the statistical significance of differences between the AUC score of the GRAPH-CNT measure and that of the other seven measures, including the baseline; the p-values are shown in Table 4 next to the corresponding AUC scores. Two observations follow from Table 4. The first is that the ordering of the measures by AUC scores is almost perfectly consistent for the two datasets. The second observation is that on the test-cro set the CNT-based measures and the baseline achieve higher AUC scores than on the test set.

Figure 1 shows the ROC curves of the best-performing coherence measures on the test dataset. As described in Section 4.2, an ROC curve measures the classification performance of a coherence measure: each point on a curve corresponds to a classifier based on a coherence measure paired with a coherence threshold used to decide whether a topic is coherent or incoherent. Figure 1 contrasts the best-performing coherence measures against the baseline measure

¹³DeLong’s test is designed to compare the AUCs of two correlated ROC curves (ROC curves derived from different measures applied to the same data points). We use the implementation from the pROC R package (Robin et al., 2011).




Figure 1: ROC curves of the best-performing coherence measures from the Table 4. The curve of the baseline measure is plotted in green on each plot. The curve of the measure from the top-performing category (top left) is plotted in red on other measures’ plots.

(green curve) and the globally best GRAPH-CNT measure (red curve). The ROC curves show that all the measures perform better than random classifiers. Furthermore, the ROC curves complement Table 4 in showing the performance gap between the GRAPH-CNT and other measures. For the GRAPH-CNT measure, the recall of 0.8 or above may be achieved with fall-out of at least 0.33. This means that if one wants a GRAPH-CNT-based classifier to detect 80% of the coherent topics, the tradeoff is to have at least 33% of the incoherent topics classified as coherent. The other measures can achieve the recall of 0.8 with fall-out rates close to 0.5 (with the exception of DENSITY-CNT yielding 0.43 fall-out). On the other hand, if one wishes to achieve fall-out rates below 0.2 (reliable detection of incoherent topics), the GRAPH-CNT measure can achieve this with a recall of 0.64 or below, while for the other measure the recall would be at best 0.56.

As seen from both Table 4 and Figure 1, the baseline measure described in Section 4.3 is a strong baseline – its performance is only slightly worse than all the proposed measures, except for the GRAPH-CNT measure, which markedly outperforms the baseline.

In the above evaluation, each of the measure categories was represented by a single best-performing measure on the development set. This rises the question whether the chosen measures are indeed representative as being the best measures within their respective categories. To answer this question, from each of the six categories we selected the top 10 or top 5% (depending on the category size) measures with the best performance on the development set (instead of a single best-performing measure per category) and evaluated these on the test and test-cro sets. Our analysis showed that the best-performing measures from Table 4 are indeed, with one exception, representative of their

Category	Parameters	AUC score	
		Dev	Test
GRAPH-CNT	<i>DocVect=tfidf, TopDocs=50, DistanceMeasure=12, DistanceThresh=0.05, Weighted=False, GraphAlgo=subgraph</i>	0.778	0.812
	<i>DocVect=tfidf, TopDocs=50, DistanceMeasure=cosine, DistanceThresh=0.02, Weighted=False, GraphAlgo=subgraph</i>	0.782	0.804
GRAPH-EMBD	<i>DocVect=cbow, TopDocs=50, DistanceMeasure=cosine, DistanceThresh=0.25, Weighted=False, GraphAlgo=subgraph</i>	0.730	0.766
	<i>DocVect=glove, TopDocs=50, DistanceMeasure=11, EmbeddingAgg=average, DistanceThresh=0.25, Weighted=True, GraphAlgo=clustering</i>	0.792	0.694
DISTANCE-CNT	<i>DocVect=bow, TopDocs=50, DistanceMeasure=cosine, DistanceAgg=average</i>	0.735	0.754
	<i>DocVect=bow, TopDocs=50, DistanceMeasure=cosine, DistanceAgg=variance</i>	0.739	0.754
DISTANCE-EMBD	<i>DocVect=cbow, TopDocs=50, DistanceMeasure=cosine, DistanceAgg=average</i>	0.711	0.746
	<i>DocVect=glove, TopDocs=25, DistanceMeasure=cosine, DistanceAgg=variance</i>	0.719	0.732
DENSITY-CNT	<i>DocVect=tfidf, TopDocs=50, DimReduce=100, CovMatrix=isotropic</i>	0.704	0.745
	<i>DocVect=tfidf, TopDocs=50, DimReduce=None, CovMatrix=isotropic</i>	0.704	0.745
DENSITY-EMBD	<i>DocVect=cbow, TopDocs=25, DimReduce=5, EmbeddingAgg=average, CovMatrix=isotropic</i>	0.701	0.734
	<i>DocVect=cbow, TopDocs=25, DimReduce=10, EmbeddingAgg=average, CovMatrix=isotropic</i>	0.708	0.728

Table 5: Parameter values of the best-performing coherence measures. For each of the six categories two best-performing measures are shown: one best performing on the development set and the other best-performing on the test set.

categories, achieving performance comparable to the measures with the best test and test-cro performance within the same category. An exception is the GRAPH-EMBD category, where the measure that performs the best on the development set does not perform well on either the test or the test-cro set. However, the selected GRAPH-EMBD measures with the top test and test-cro performances perform markedly better. A detailed analysis of the representativeness of the best development set measures is provided in the supplementary material.¹⁴

We now turn to a question of practical importance: which parameter values yield the best measures, i.e., which measures should one apply in practice? In Table 5, we give an indicative summary which shows the parameter values for

¹⁴<https://rebrand.ly/doc-coh-supplementary>

two best-performing measures from each of the categories – one with the best performance on the development set and another with the best performance on the test set. A detailed analysis of the parameters of the best measures is provided in the supplementary material,¹⁴ while here we summarize only the most interesting findings. Interestingly, all of the best-performing GRAPH-CNT measures use thresholding to filter the edges, but ultimately use an unweighted graph. Of the five proposed graph-based coherence algorithms (cf. Section 3.3.3), three algorithms emerge as components of the best-performing measures: **subgraph-centrality**, **closeness-centrality**, and **clustering-coefficient**. Common to all three algorithms is that they calculate local connectivity scores of graph nodes and average them to obtain the graph score, unlike the other two algorithms (**num-connected** and **min-spanning-tree**), which calculate global graph connectivity properties. As regards the edge threshold value, **subgraph** and **closeness** centrality algorithms generally prefer smaller thresholds yielding sparse graphs (percentile ranks 0.02, 0.05, and 0.10), while the **clustering** algorithm seems to prefer higher thresholds (percentile ranks 0.25 and 0.5). The best GRAPH-EMBD measures have the same structure, but fail to achieve the performance of the best GRAPH-CNT measures. All the best measures that use EMBD-based document vectorization in combination with **l1** and **l2** distances perform the aggregation of document vectors by averaging instead of summing the word embeddings. This is expected because the **l1** and **l2** distance measures, unlike the cosine distance, are not invariant to vector length. Generally, representing documents with CNT vectors seems preferable over EMBD vectors. More specifically, for graph-based measures the choice of CNT vectors over EMBD vectors leads to a large increase in performance, yielding globally best performance.

4.6. Conclusions

The main finding from the above experiments is that the best overall performance is achieved by coherence measures from the GRAPH-CNT category, i.e., measures that rely on count-based document vectorization derived from an in-domain corpus and combined with graph-based algorithms for coherence scoring. This result was confirmed on both the US news topics dataset and the Croatian news topics dataset. Measures from the GRAPH-CNT category construct graphs of top topic documents by removing all edges above a small distance threshold and calculate the coherence score by aggregating local node connectivity information, suggesting that the most effective way to estimate coherence of a set of documents is to use vectors of word-in-document counts to represent the documents and average local similarities in the neighborhoods of each document. In contrast, the measures from the GRAPH-EMBD category, which fail to reach the performance of best GRAPH-CNT measures, also average local similarities but use embedding-based document vectorization. Measures from other four non-GRAPH categories calculate the global similarity of the entire set of documents.

Another observation arising from the experiments is that, not surprisingly, the choice of document vectorization method has a strong influence on the per-

formance of the coherence measures. These observations indicate that improvements to existing methods could be achieved by means of some other vectorization method.

Finally, the choice of top 50 topic-related documents as input for calculation of a topic’s coherence score emerged as the best option that results in measures achieving best or nearly-best results for all the measure categories, as demonstrated by the results in Section 4.5 (and additional analyses provided the supplementary material).

5. Experiments with Word-based Coherence

The previous experiment has investigated the efficacy of document-based coherence measures. In this section, we turn to the question of how document-based coherence measures, which we propose in this work, relate to the word-based coherence. While in the introduction we argued that document-based coherence measures may be better suited for topics derived from news media texts, the objective of the experiments described in this section is to quantify the magnitude and consistency of these gains. To this end, we measure how well the state-of-art word-based coherence measures approximate document-based coherence of topics derived from news media texts. Lastly, to gain additional insight into differences between the two types of measures, we perform a qualitative analysis of topics coming from the high and low ends of the document- and word-based coherence spectrum.

5.1. Word-Coherence Measures

As a reference point for selecting the word-based coherence measures, we use the study of Röder et al. (2015), who carried out a detailed and systematic analysis of a large number of word-based coherence measures on six publicly available datasets. For our experiments, we select five top-performing measures from this study, designed to predict coherence scores formed by inspection of top-ranked topic words: (1) the C_{UCI} measure of Newman et al. (2010), (2) the C_{NPMI} measure of Aletras & Stevenson (2013), (3) the C_A measure of Aletras & Stevenson (2013), and the (4) C_V and (5) C_P measures, both discovered by an exhaustive search of the parameter space in (Röder et al., 2015).¹⁵

In the generic framework of Röder et al. (2015), the above measures are defined by an appropriate partitioning of the set of topic words into subsets, followed by computing the average of an appropriate similarity measure¹⁶ between all pairs of so-obtained subsets. This includes the calculation of similarity of word pairs as a special case. The pairwise similarities between word subsets are calculated from probabilities based on word co-occurrences derived from a

¹⁵All five considered measures are implemented in the open-source software package Palmetto, available at <https://github.com/dice-group/Palmetto>.

¹⁶Röder et al. (2015) use the term “confirmation measure” to refer to a similarity measure between subsets of topic words.

corpus, which are affected by the choice of the corpus and the choice of the corpus preprocessing method (tokenization, stopword removal, and word normalization).

The measures C_{UCI} and C_{NPMI} average the similarity of word pairs computed using pointwise mutual information (PMI) and its normalized version (NPMI), respectively. A sliding window is used to derive word co-occurrence for both measures. The C_A measure also averages the similarity of word pairs, with the difference that the words in the pair are first represented as vectors of NPMI-based similarities with other top-ranked topic words and the similarity of the pair is calculated as the similarity of these two vectors. The C_V measure averages the similarities between pairs each comprising a topic word and its complement set that consists of all other topic words. Similarly, as for C_A , indirect vector similarity is used to compare a word with its complement set. Lastly, the C_P measure averages similarities of pairs each comprising a topic word and all topic words ranked above it by the word-topic probabilities, using as similarity a measure based on conditional probability of a single word given a word set (Fitelson, 2003).

5.2. Estimating Document-Based Coherence with Word-Based Coherence

In this experiment, we examine how well the top-performing word-based coherence measures approximate document-based coherence scores. Namely, word-based coherence measures from Section 5.1 are designed to predict, using as input a set of top-ranked topic words, coherence scores of topics assigned by human annotators based on the inspection of top-ranked topic *words*. In contrast, document-based coherence scores are assigned by annotators who inspected the top-ranked topic *documents*, and accordingly document-based measures considered in Section 4 use as input the top-ranked topic documents. We evaluate word-based coherence measures in the same way as we have evaluated the document-based measures, namely, on the test and test-cro sets of manually annotated topics using AUC score as the performance measure (cf. Sections 4.1 and 4.2).

We use the measures with the parameters optimized for word-based coherence estimation (Röder et al., 2015), all of which use top-10 topic words as input and derive the co-occurrence counts from Wikipedia. For the US news topics dataset, we derive the counts from the English Wikipedia (dump from June 2016), while for the Croatian news topics dataset, we derive the counts from the Croatian Wikipedia (dump from November 2017). In both cases the counts are derived with the preprocessing we used when we built our topic models.¹⁷ Before preprocessing, we removed the redirection, disambiguation, category, and portal pages from both Wikipedia datasets.

¹⁷ For the US topics we also tried using the original counts used by Röder et al. (2015), available online, but the counts obtained with our preprocessing turned out to give better AUC scores for all the measures.

Measure	test		test-cro	
	AUC	p-value	AUC	p-value
<i>doc-dist-cosine</i>	0.730	–	0.748	–
C_V	0.607	0.002	0.508	0.000
C_A	0.579	0.001	0.442	0.000
C_P	0.548	0.000	0.614	0.009
C_{NPMI}	0.498	0.000	0.595	0.002
C_{UCI}	0.482	0.000	0.571	0.001

Table 6: Performance of word-based coherence methods in estimating document-based coherence, compared against the document-based baseline coherence measure.

Table 6 shows the performance of state-of-art word-based coherence measures compared against the baseline document-based coherence method *doc-dist-cosine* (cf. Section 4.3). The accompanying p-values are obtained using the DeLonge’s test (cf. Section 4.5) with the null hypothesis of no difference between the AUC scores of a word-based measure and that of the baseline. As is evident from the results, document-based coherence measures outperform word-based measures by a considerable margin – the best word-based measures achieve AUC scores slightly above 0.6, while the document-based baseline achieves scores of at least 0.73. Partitioning the set of top-ranked topic words into pairs of words and word sets, as implemented by C_V and C_P measures, seems to yield somewhat better AUC scores than partitioning into word pairs, as implemented by all other word-based coherence measures.

5.3. Qualitative Analysis

The previous experiment has shown that, when comparing to the proposed document-based coherence measures, state-of-art word-based coherence measures fall short of estimating document-based coherence. This raises the question of what is the relation between these two approaches to measuring coherence: is document-based coherence simply a better model of topic coherence, or are word- and document-based coherence two different, although correlated and possibly complementary views on topic coherence?

To investigate which of the two is the case, we carried out a qualitative analysis of the topics from our US news topics dataset. The analysis is done along two dimensions: document-based coherence and word-based coherence. In each of the two dimensions, we select topics of high coherence and topics of low coherence, giving us four categories of topic coherence. The topics are selected from a sample of 230 topics constituting the test set (cf. Section 4.1), based on scores produced by the document- and word-based coherence measures. More specifically, for document-based coherence we select from the top 30% and bottom 30% of topics ranked by the best-performing measure from the GRAPH-CNT category (cf. Section 4.5), while for word-based coherence we do the same using the C_P measure (this measure performs well on word-based coherence

Topic label	Top-10 topic words
Chicago mayoral election	mayor chicago emanuel de giuliani garcia love blasio rudy runoff
Restoration of ties with Cuba	foundation cuba list malley summit rubio cuban donations trump island
Ted Cruz	cruz ted liberty tea imagine evangelical r-texas declared candidacy obamacare
Iran negotiations	nuclear agreement sanctions iranian weapons kerry framework tehran cotton corker
Vaccination	vaccines parents science kids choice huffpost carson measles research believes

Table 7: Topics with high document coherence and low word coherence.

but poorly on document-based coherence; cf. Table 6). Furthermore, for the purpose of this analysis, we manually categorize each semantic topic as either *concrete* (topics pertaining to entities, events, or news stories) or *abstract* (topics pertaining to general issues or abstract semantic categories).

High document/low word coherence. In this category, we find 23 topics from the test set, the majority (21) of which are concrete. Table 7 gives five examples, where all but the “vaccination” topic are concrete. The low word-based coherence scores are in line with the observation that the top-ranked topic words are, as a whole, semantically unrelated. However, high document-based coherence scores are a consequence of the topic-related documents being highly similar news articles describing the same entity, event, or story. Note that, since a state-of-art word-based coherence measure is unable to recognize the coherence of these topics, they would likely be discarded as low-quality topics. Likewise, a human annotator – unless familiar with the news corpus – would judge the set of top-ranked words as unrelated and the topics as incoherent. Hence, high document/low word coherence topics are a paradigmatic case for the use of document-based coherence measures.

High document/high word coherence. Among the 20 topics in this category, 14 are abstract and 6 are concrete. This suggests that topics scored with high word-based coherence tend to be abstract. The top-ranked topic words of such abstract topics are semantically related. Table 8 shows five example topics from this category. Among these, two are concrete topics pertaining to entities, which shows that concrete topics can also have high word coherence if they are characterized by words relating to a common concept. On the other hand, the fact that 14 abstract topics are also scored as coherent by the document-based coherence measure suggests that document-based coherence can detect the coherence of not only concrete but also abstract semantic topics.

Topic label	Top-10 topic words
Environment	climate energy global science environmental warming fuel scientists emissions plants
Budget	billion domestic fiscal balance deficit medicare repeal priorities ryan trillion
Consumer debt crisis	debt loans dollars contract fees payments taxpayers borrowers treasury consumers
Robert Menendez	attorney menendez lawyer criminal allegations file sentence prosecutors convicted prison
Yemen	saudi strike target al yemen intelligence arabia houthis pakistan qaeda

Table 8: Topics with high document coherence and high word coherence.

Low document/high word coherence. From the 12 topics in this category, three are incoherent, while the rest are abstract topics. This again confirms the correlation between topic abstractness and word-based coherence. The abstract topics fall into two groups. The first is a group of four topics (lawsuits, journalism, social media, and radio & television) whose low document coherence is due to topics being unrepresented in the documents, i.e., they are mentioned in a relatively small portion of document text otherwise dominated by other topics. This makes the topic-related documents heterogeneous and incoherent, which in turn lowers the document-based coherence score. The second group are five topics that are coherent but for which the document-based coherence score is either misestimated or low because a topic is highly abstract and associated with a set of documents that are less semantically related. Table 9 shows all the topics from the first group and one example topic from the second group. Taken together, this category of topics shows that it might be beneficial to combine document- and word-based coherence measures: the word coherence could be used as a fallback in cases when a document-based coherence measure fails to detect coherence, for instance because the topic is being underrepresented in the documents.

Low document/low word coherence. Among the 26 topics in this category, the majority of topics (18) are incoherent (mixture of topics or noise), while the remaining 8 topics pertain to a single semantic topic but with the addition of noise. Additionally, the majority of the remaining topics (7 out of 8) are concrete – a property that correlates with low word-based coherence. The remaining set of 8 topics in this category represents a hard case whose coherence is difficult to detect even with a combination of a document- and word-based coherence measure – these topics are specific in that they contain both noise (which decreases document-based coherence) and are concrete (which decreases word-based coherence). This category of topics complements the one containing topics with low document and high word coherence in demonstrating the useful-

Topic label	Top-10 topic words
Lawsuits	file board lawsuit complaint suit violated georgia damage settled accused
Radio and television	morning host watch night radio television tv network update station
Journalism	writing published article piece journalists paper journal newspaper quoted editor
Social media	fox twitter host night tweeted morning facebook com watch remarks
Crime	prison criminal crime sentence convicted attorney prosecutors trial lawyer judge
<i>(noise)</i>	video someone thought probably maybe else guy anything everyone yes

Table 9: Topics with low document coherence and high word coherence.

ness of combining word-based and document-based coherence measures. Among the topics with low document coherence, word coherence correlates with true coherence: most (18 out of 26) low word coherence topics are indeed incoherent, while most (10 out of 12) high word coherence topics are indeed coherent.

5.4. Conclusions

We motivated the need for document-based coherence by essentially claiming that word-based coherence in some cases is not informative enough to gauge the coherence of a topic derived by a topic model, especially for transient and contingent topics typical of the news domain. Results of the experiments in Section 5.2 confirm this by showing that when state-of-art word-based coherence measures proposed in the literature are used for estimating document-based coherence, their performance is markedly below the document-based coherence baseline. However, some word-based measures outperform a random baseline, indicating a degree of correlation between document-based and word-based coherence.

The qualitative analysis of model topics with high and low document- and word-based coherence indicates that word- and document-based coherence measures are complementary to each other and that they may be combined to detect coherent topics more accurately – there exist coherent topics (model topics corresponding to semantic topics) that would be discarded as incoherent if the detection were based on word- or document-based coherence alone. Interestingly, high word-based coherence correlates with abstract topics, whereas low word-based coherence correlates with concrete topics.

6. Semi-Automated Topic Discovery

As mentioned in the introduction, the increased availability and consumption of online news in textual form brought about increased interest in automated content analysis of news texts. Most content analysis techniques rely on an inventory of topics, which ideally corresponds to the topics covered by the news texts under analysis. However, large quantities of news data available on one hand, and the transience of topics reported about in the news on the other, make it impossible to establish a fixed and comprehensive topic inventory for the news domain. The alternative is to rely on techniques for (semi)automated topic discovery from text corpus, for which topic models have proven to be an extremely useful tool.

However, although topic models make topic discovery much easier, for the reasons discussed in the introduction the results are of varying quality – more precisely, not all topics induced by a topic model will be semantically interpretable. Hence, to be usable for content analysis, the results of the topic model typically need to be examined by a human expert, which requires considerable effort. Moreover, to increase the coverage of the topics, one would typically want to run several differently parametrized topic models (e.g., models with varying number of topics) and examine the topics collected from all these models, which further increases the complexity of the task.

In this section we demonstrate how our proposed document-based coherence measures can be used to improve the efficiency of semi-automated topic discovery based on topics collected from several models. The idea is to use document-based topic coherence scores as a heuristic for which topics should be examined first by the human expert. The assumption is that, if the human expert examines the topics in the order of their coherence score rather than at random, and if care is taken to avoid re-discovering duplicate topics, this may substantially improve the discovery rate of the semantic topics (the number of discovered topics per number of topics examined), thus reducing the overall human effort.

6.1. *Experimental Setup*

We simulate the described topic discovery scenario using our US news topics dataset of model topics annotated with semantic topics (cf. Section 4.1) by traversing the topics sorted in descending order of coherence score and keeping a count of the number of distinct semantic topics.

We simulate the semantic topic discovery process using the entire US news topics dataset consisting of 350 model topics. Each coherent topic is labeled with a single corresponding semantic topic, while both fused topics and noise topics are labeled with zero semantic topics. Discovery is simulated by traversing the list of model topics sorted either at random or by scores assigned by a coherence measure. Each model topic is compared with the topics from the list of already discovered topics and discarded if it is found to be a duplicate. If the topic is not a duplicate, it is counted as a topic examined by a human expert, and if the topic corresponds to a semantic topic, the semantic topic is added to the list

of discovered topics. We measure the effectiveness of the discovery process by examining the topic discovery rate: the number of discovered semantic topics per number of topics examined.

Two points deserve to be mentioned: topic duplication and fused topics. While using more than a single topic model will generally improve the coverage of the semantic topics, it will also result in some topics being duplicated, especially those that are more salient in the corpus. Because incoherent topics are random, being the result of either random noise or of the fusing of two random semantic topics, they are very unlikely to be duplicated. Alternately, coherent topics will match semantic topics, which are limited in number, and therefore topics that get duplicated are most likely coherent. As a consequence, sorting the topics by coherence will push the duplicates toward the start of the list, thereby lowering the topic discovery rate, as each duplicate topic needs to be examined but yields no new semantic topics. To prevent this, we remove all duplicates from the pooled list of topics. We consider two topics to be duplicates if the cosine distance between the topics’ probability vectors is less than 0.5 – a conservative choice, as topics have to be almost identical to meet this threshold. When simulating the topic discovery process, we apply duplicate removal to both topics sorted by coherence and randomly ordered topics.

Another point worth mentioning concerns the fused topics: model topics corresponding to two or more semantic topics. As described in Section 4.1, 21% of model topics in our dataset were labeled with two semantic topics. In this experiment, we treat fused topics as noise, simulating the topic discovery scenario in which the fused topics are recognized by the human expert as noise and discarded – a task that can be performed efficiently. Some semantic topics that only occur as fused may be missed, but using a pool of topic models rather than a single model makes it more likely that the fused topic will eventually emerge as separate and coherent topic in one of the models. Alternatively, the human expert could attempt to split up fused topics into individual semantic topics but this would be a time-consuming and potentially error-prone task. In this approach to topic discovery the benefit of using a coherence measure to improve the discovery rate would be reduced, since the fused topics generally have low coherence. To remedy this, the coherence measures could be compounded with measures for detection of fused topics, but we leave this for future work.

6.2. Results

We tested four coherence measures: the baseline document-based measure (cf. Section 4.3), the C_P word-based measure (cf. Section 5.1), and two top-performing document-based measures from the GRAPH-CNT category with the highest AUC score on the development set and the test set (cf. Table 5), subsequently denoted **graph1** and **graph2**, respectively. We compare the four coherence measures against a random baseline: we run the simulation for 20 random topic orderings and calculate the average, minimum, and maximum of discovered topics per the number of examined topics.

The results are shown in Figure 2. The curves on the left plot show the number of topics discovered as a function of the number of topics examined,




Figure 2: Semantic topics discovered (y-axis) per model topics examined (x-axis): (a) absolute number of topics discovered and (b) the average difference between the number of topics discovered when using ordering based on a topic coherence measure and random ordering.

while the right plot shows the difference between the number of topics discovered and the average number of topics discovered when examined in random order. The main observation is that topic discovery guided by the document-based coherence measures markedly outperforms both the random baseline and the state-of-art word-based coherence measure. Furthermore, the two GRAPH-CNT coherence measures outperform the document-based coherence baseline, giving the overall best performance. Another observation is that these two coherence measures, albeit they score nearly identical in terms of AUC (cf. Table 5), yield different topic discovery curves. This suggests that for evaluating coherence measures for specific applications it might be a good idea to complement rank-based metrics such as AUC with an application-oriented evaluation metric.

In more concrete terms, the results on this dataset show that by sorting the topics by coherence-based coherence it is possible to discover all semantic topics after examining 160 model topics – this is in contrast to an average of 200 model topics which would have to be examined if the order were random. Assuming that topic examination takes 6 minutes on average (an estimate based on the annotation of the dataset), the coherence measure would save the annotator four hours. If the aim is to examine only 100 model topics, perhaps to get an overview of semantic topics covered by the corpus, examination in random order would on average discover only 59 semantic topics, whereas examination based on document-coherence would discover 77 topics – an increase of 30%.

7. Conclusion

Topic models are a popular tool for unsupervised discovery of topics from text corpora, including texts from the news domain. A well-known problem with topic models, however, is that the quality of the generated topics typically varies. This motivated the development of a number of model evaluation techniques, most notably those based on the calculation of topics’ semantic coherence. The existing topic coherence methods estimate the coherence based on the semantic relatedness of the topic words. This approach, however, is inadequate for news

media texts, where topics are often contingent and transient and therefore associated with semantically unrelated words. To solve this problem, we proposed a novel class of topic coherence methods that estimate topic coherence based on topic documents rather than topic words. The underlying assumption is that, because documents contain more information than words, document-based topic coherence can better capture topics' semantic interpretability.

The proposed document-based methods calculate the coherence of a topic in three steps: selection of topic-related documents, vectorization of the documents, and calculation of a coherence score from document vectors. We proposed a number of different methods, including distance-based, density-based, and graph-based methods, and evaluated them on two datasets of topics manually labeled with coherence scores. The method that uses tf-idf or bag-of-words document representations, builds an unweighted similarity graph, and estimates the coherence score by aggregating node connectivity scores was found to outperform all other considered methods in terms of coherence ranking performance, including a strong baseline. Furthermore, we have shown that the method can be used to speed up the otherwise tedious task of semi-automated topic discovery from a corpus of news media texts.

To investigate the relationship between document-based and word-based topic coherence, we evaluated state-of-art word-based coherence methods on our datasets of topics labeled with document-based coherence scores. We found that word-based coherence methods, which are optimized for word-based coherence, perform poorly on our datasets. An examination of model topics with estimated high and low document- and word-based coherence demonstrated the potential merit in combining word- and document-based coherence measures to detect coherent topics more accurately.

There exist a number of interesting directions for future work. On the technical side, the methods we propose could probably be improved by using a more effective document vectorization method. The experiments in Section 4.5 show that the baseline document-based method exhibits fairly good performance, so one possibility might be to try to improve the baseline or combine it with the measures we proposed. Another promising direction for future work, indicated by the experiments in Section 5, would be to combine document- and word-based coherence measures.

At a conceptual level, we believe that further investigating the possible applications of both word- and document-based coherence measures in exploratory analysis might prove these measures useful for a task separate from topic model evaluation. Such applications might also help us gain a better understanding of the concept of topic coherence. Document-based coherence measures might be used for a systematic comparative analysis of the different existing topic models variants. Ideally, such an analysis would also consider the application of document-based coherence to other genres of texts besides news texts, such as the topically more heterogeneous scientific articles.

8. Acknowledgments

This research has been supported by the European Regional Development Fund under the grant KK.01.1.1.01.0009 (DATACROSS). We would like to thank Domagoj Alagić, Mladen Karan and Maria Pia Di Buono for the helpful feedback.

References

- Ahmed, A., Ho, Q., Eisenstein, J., Xing, E., Smola, A. J., & Teo, C. H. (2011). Unified analysis of streaming news. In *Proceedings of the 20th international conference on World wide web* (pp. 267–276). ACM.
- Aletras, N., & Stevenson, M. (2013). Evaluating topic coherence using distributional semantics. In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013)* (pp. 13–22).
- Allan, J. (2012). *Topic detection and tracking: event-based information organization* volume 12. Springer Science & Business Media.
- AlSumait, L., Barbará, D., Gentle, J., & Domeniconi, C. (2009). Topic significance ranking of lda generative models. *Machine Learning and Knowledge Discovery in Databases*, (pp. 67–82).
- American Press Institute (2015). How millennials get news: Inside the habits of america’s first digital generation.
- Balahur, A., Steinberger, R., Kabadjov, M., Zavarella, V., Van Der Goot, E., Halkia, M., Pouliquen, B., & Belyaeva, J. (2013). Sentiment analysis in the news. *arXiv preprint arXiv:1309.6202*, .
- Belford, M., Namee, B. M., & Greene, D. (2018). Stability of topic modeling via matrix factorization. *Expert Systems with Applications*, *91*, 159 – 169. doi:<https://doi.org/10.1016/j.eswa.2017.08.047>.
- Blei, D. M. (2012). Probabilistic topic models. *Communications of the ACM*, *55*, 77–84.
- Blei, D. M., & Lafferty, J. D. (2007). A correlated topic model of science. *The Annals of Applied Statistics*, *1*, 17–35.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, *3*, 993–1022.
- Bogers, T., & Van den Bosch, A. (2007). Comparing and evaluating information retrieval algorithms for news recommendation. In *Proceedings of the 2007 ACM conference on Recommender systems* (pp. 141–144). ACM.

- Boyd-Graber, J. L., Blei, D. M., & Zhu, J. (2007a). Probabilistic walks in semantic hierarchies as a topic model for WSD. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- Boyd-Graber, J. L., Blei, D. M., & Zhu, X. (2007b). A topic model for word sense disambiguation. In *EMNLP-CoNLL* (pp. 1024–1033).
- Chang, J., Boyd-Graber, J. L., Gerrish, S., Wang, C., & Blei, D. M. (2009). Reading tea leaves: How humans interpret topic models. In *Nips* (pp. 1–9). volume 31.
- Chen, M., Jin, X., & Shen, D. (2011). Short text classification improved by learning multi-granularity topics. In *IJCAI* (pp. 1776–1781).
- Chuang, J., Fish, S., Larochelle, D., Li, W. P., & Weiss, R. (2014). Large-scale topical analysis of multiple online news sources with media cloud. *NewsKDD: Data Science for News Publishing, at KDD*, .
- Chuang, J., Gupta, S., Manning, C., & Heer, J. (2013). Topic model diagnostics: Assessing domain relevance via topical alignment. In *Proceedings of the 30th International Conference on machine learning (ICML-13)* (pp. 612–620).
- Chuang, J., Ramage, D., Manning, C., & Heer, J. (2012). Interpretation and trust: Designing model-driven visualizations for text analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 443–452). ACM.
- Clerwall, C. (2014). Enter the robot journalist: Users’ perceptions of automated content. *Journalism Practice*, 8, 519–531.
- Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American society for information science*, 41, 391.
- DeLong, E. R., DeLong, D. M., & Clarke-Pearson, D. L. (1988). Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics*, (pp. 837–845).
- Dou, W., Wang, X., Skau, D., Ribarsky, W., & Zhou, M. X. (2012). Leadline: Interactive visual analysis of text data through event identification and exploration. In *Visual Analytics Science and Technology (VAST), 2012 IEEE Conference on* (pp. 93–102). IEEE.
- Entman, R. M. (1993). Framing: Toward clarification of a fractured paradigm. *Journal of communication*, 43, 51–58.
- Estrada, E., & Rodriguez-Velazquez, J. A. (2005). Subgraph centrality in complex networks. *Physical Review E*, 71, 056103.
- Fitelson, B. (2003). A probabilistic theory of coherence. *Analysis*, 63, 194–199.

- Flaounas, I., Ali, O., Lansdall-Welfare, T., De Bie, T., Mosdell, N., Lewis, J., & Cristianini, N. (2013). Research methods in the age of digital journalism: Massive-scale automated analysis of news-content topics, style and gender. *Digital Journalism*, 1, 102–116.
- Freeman, L. C. (1978). Centrality in social networks conceptual clarification. *Social networks*, 1, 215–239.
- Gao, W., Li, P., & Darwish, K. (2012). Joint topic modeling for event summarization across news and social media streams. In *Proceedings of the 21st ACM international conference on Information and knowledge management* (pp. 1173–1182). ACM.
- Garcin, F., Dimitrakakis, C., & Faltings, B. (2013). Personalized news recommendation with context trees. In *Proceedings of the 7th ACM conference on Recommender systems* (pp. 105–112). ACM.
- Godbole, N., Srinivasaiah, M., & Skiena, S. (2007). Large-scale sentiment analysis for news and blogs. *ICWSM*, 7, 219–222.
- Griffiths, T. L., & Steyvers, M. (2004). Finding scientific topics. *Proceedings of the National academy of Sciences*, 101, 5228–5235.
- Grimmer, J. (2009). A bayesian hierarchical topic model for political texts: Measuring expressed agendas in senate press releases. *Political Analysis*, 18, 1–35.
- Grimmer, J., & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, (pp. 1–31).
- Grobelnik, M., & Mladenic, D. (2004). Visualization of news articles. *Informatika (Slovenia)*, 28, 375–380.
- Hoffman, M., Bach, F. R., & Blei, D. M. (2010). Online learning for latent Dirichlet allocation. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, & A. Culotta (Eds.), *Advances in Neural Information Processing Systems 23* (pp. 856–864). Curran Associates, Inc.
- Jacobi, C., van Atteveldt, W., & Welbers, K. (2016). Quantitative analysis of large amounts of journalistic texts using topic modelling. *Digital Journalism*, 4, 89–106.
- Kim, Y., Kim, S., Jaimes, A., & Oh, A. (2014). A computational analysis of agenda setting. In *Proceedings of the Companion Publication of the 23rd International Conference on World Wide Web Companion* (pp. 323–324). ACM.
- Koltcov, S., Koltsova, O., & Nikolenko, S. (2014). Latent dirichlet allocation: stability and applications to studies of user-generated content. In *Proceedings of the 2014 ACM conference on Web science* (pp. 161–165). ACM.

- Korenčić, D., Ristov, S., & Šnajder, J. (2015). Getting the agenda right: measuring media agenda using topic models. In *Proceedings of the 2015 Workshop on Topic Models: Post-Processing and Applications* (pp. 61–66). ACM.
- Korenčić, D., Grbeša-Zenzerović, M., & Šnajder, J. (2016). Topics and their salience in the 2015 parliamentary election in Croatia: A topic model based analysis of the media agenda. In *Proceedings of the International Conference on the Advances in Computational Analysis of Political Text - PolText 2016*.
- Lau, J. H., Newman, D., & Baldwin, T. (2014). Machine reading tea leaves: Automatically evaluating topic coherence and topic model quality. In *EACL* (pp. 530–539).
- Lee, C.-S., Chen, Y.-J., & Jian, Z.-W. (2003). Ontology-based fuzzy event extraction agent for chinese e-news summarization. *Expert Systems with Applications*, 25, 431–447.
- Lee, D. D., & Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, 401, 788.
- Li, L., Wang, D.-D., Zhu, S.-Z., & Li, T. (2011). Personalized news recommendation: a review and an experimental investigation. *Journal of computer science and technology*, 26, 754–766.
- Li, L., Zheng, L., Yang, F., & Li, T. (2014). Modeling and broadening temporal user interest in personalized news recommendation. *Expert Systems with Applications*, 41, 3168–3177.
- Lin, C., & He, Y. (2009). Joint sentiment/topic model for sentiment analysis. In *Proceedings of the 18th ACM conference on Information and knowledge management* (pp. 375–384). ACM.
- Ling, C. X., Huang, J., & Zhang, H. (2003). Auc: a statistically consistent and more discriminating measure than accuracy. In *IJCAI* (pp. 519–524). volume 3.
- Lippmann, W. (1922). *Public Opinion*. Macmillan.
- Ljubešić, N., & Erjavec, T. (2011). hrWaC and slWac: Compiling web corpora for Croatian and Slovene. In *Text, Speech and Dialogue - 14th International Conference Lecture Notes in Computer Science* (pp. 395–402). Springer.
- Magnini, B., Strapparava, C., Pezzulo, G., & Gliozzo, A. (2002). The role of domain information in word sense disambiguation. *Natural Language Engineering*, 8, 359–373.
- Masand, B., Linoff, G., & Waltz, D. (1992). Classifying news stories using memory based reasoning. In *Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval* (pp. 59–65). ACM.

- McCombs, M. (2014). *Setting the Agenda: The Mass Media and Public Opinion*. Polity Press.
- McCombs, M. E., & Shaw, D. L. (1972). The agenda-setting function of mass media. *Public Opinion Quarterly*, 36, 176–187.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781. URL: <http://arxiv.org/abs/1301.3781>. arXiv:1301.3781.
- Mimno, D., & Blei, D. (2011). Bayesian checking for topic models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing* (pp. 227–237). Association for Computational Linguistics.
- Mimno, D., Wallach, H. M., Talley, E., Leenders, M., & McCallum, A. (2011). Optimizing semantic coherence in topic models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing* (pp. 262–272). Association for Computational Linguistics.
- Mimno, D. M., & McCallum, A. (2012). Topic models conditioned on arbitrary features with Dirichlet-multinomial regression. *CoRR*, abs/1206.3278.
- Musat, C., Velcin, J., Trausan-Matu, S., & Rizoïu, M.-A. (2011). Improving topic evaluation using conceptual knowledge. In *22nd International Joint Conference on Artificial Intelligence (IJCAI)* (pp. 1866–1871). volume 3.
- Nassirtoussi, A. K., Aghabozorgi, S., Wah, T. Y., & Ngo, D. C. L. (2015). Text mining of news-headlines for forex market prediction: A multi-layer dimension reduction algorithm with semantics and sentiment. *Expert Systems with Applications*, 42, 306–324.
- Neuendorf, K. A. (2016). *The content analysis guidebook*. Sage.
- Newman, D., Chemudugunta, C., Smyth, P., & Steyvers, M. (2006). Analyzing entities and topics in news articles using statistical topic models. In *ISI* (pp. 93–104). Springer.
- Newman, D., Lau, J. H., Grieser, K., & Baldwin, T. (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* (pp. 100–108). Association for Computational Linguistics.
- Newman, N., Fletcher, R., Kalogeropoulos, A., Levy, D., & Kleis Nielsen, R. (2016). Reuters institute digital news report 2017.
- Nikolenko, S. I. (2016). Topic quality metrics based on distributed word representations. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval* (pp. 1029–1032). ACM.

- Nikolenko, S. I., Koltcov, S., & Koltsova, O. (2015). Topic modelling for qualitative studies. *Journal of Information Science*, *43*, 88–102.
- O’Callaghan, D., Greene, D., Carthy, J., & Cunningham, P. (2015). An analysis of the coherence of descriptors in topic modeling. *Expert Systems with Applications*, *42*, 5645–5657.
- Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. In *EMNLP* (pp. 1532–1543). volume 14.
- Phuvipadawat, S., & Murata, T. (2010). Breaking news detection and tracking in twitter. In *Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on* (pp. 120–123). IEEE volume 3.
- Piskorski, J., & Yangarber, R. (2013). Information extraction: past, present and future. In *Multi-source, multilingual information extraction and summarization* (pp. 23–49). Springer.
- Popescu, O., & Strapparava, C. (Eds.) (2017). *Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism*. Copenhagen, Denmark: Association for Computational Linguistics. URL: <http://www.aclweb.org/anthology/W17-42>.
- Ramirez, E. H., Brena, R., Magatti, D., & Stella, F. (2012). Topic model validation. *Neurocomputing*, *76*, 125–133.
- Ramrakhiani, N., Pawar, S., Hingmire, S., & Palshikar, G. K. (2017). Measuring topic coherence through optimal word buckets. *EACL 2017*, (p. 437).
- Roberts, M., Stewart, B., & Tingley, D. (2016). Navigating the local modes of big data: The case of topic models. In *Computational Social Science* (pp. 51–97). Cambridge University Press, New York.
- Robin, X., Turck, N., Hainard, A., Tiberti, N., Lisacek, F., Sanchez, J.-C., & Müller, M. (2011). pROC: an open-source package for r and s+ to analyze and compare roc curves. *BMC bioinformatics*, *12*, 77.
- Röder, M., Both, A., & Hinneburg, A. (2015). Exploring the space of topic coherence measures. In *Proceedings of the eighth ACM international conference on Web search and data mining* (pp. 399–408). ACM.
- Rose, T., Stevenson, M., & Whitehead, M. (2002). The reuters corpus volume 1 - from yesterday’s news to tomorrow’s language resources. In *LREC* (pp. 827–832). volume 2.
- Rosner, F., Hinneburg, A., Röder, M., Nettling, M., & Both, A. (2014). Evaluating topic coherence measures. *CoRR*, *abs/1403.6397*. URL: <http://arxiv.org/abs/1403.6397>. arXiv:1403.6397.

- Salton, G., & Buckley, C. (1988). Term-weighting approaches in automatic text retrieval. *Information processing & management*, 24, 513–523.
- Saramäki, J., Kivelä, M., Onnela, J.-P., Kaski, K., & Kertesz, J. (2007). Generalizations of the clustering coefficient to weighted complex networks. *Physical Review E*, 75.
- Schult, D. A., & Swart, P. (2008). Exploring network structure, dynamics, and function using networkx. In *Proceedings of the 7th Python in Science Conferences (SciPy 2008)* (pp. 11–16). volume 2008.
- Shahaf, D., & Guestrin, C. (2012). Connecting two (or less) dots: Discovering structure in news articles. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 5, 24.
- Steinberger, R., Pouliquen, B., & der Goot, E. V. (2013). An introduction to the europe media monitor family of applications. *CoRR*, abs/1309.5290. URL: <http://arxiv.org/abs/1309.5290>. arXiv:1309.5290.
- Stevens, K., Kegelmeyer, P., Andrzejewski, D., & Buttler, D. (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* (pp. 952–961). Association for Computational Linguistics.
- Tanev, H., Piskorski, J., & Atkinson, M. (2008). Real-time news event extraction for global crisis monitoring. *Natural Language and Information Systems*, (pp. 207–218).
- Titov, I., & McDonald, R. T. (2008). A joint model of text and aspect ratings for sentiment summarization. In *ACL* (pp. 308–316).
- Turian, J., Ratinov, L., & Bengio, Y. (2010). Word representations: a simple and general method for semi-supervised learning. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics* (pp. 384–394). Association for Computational Linguistics.
- Vossen, P., Rigau, G., Serafini, L., Stouten, P., Irving, F., & Van Hage, W. R. (2014). NewsReader: recording history from daily news streams. In *LREC* (pp. 2000–2007).
- Waal, A. D., & Barnard, E. (2008). Evaluating topic models with stability.
- Wallach, H. M., Murray, I., Salakhutdinov, R., & Mimno, D. (2009). Evaluation methods for topic models. In *Proceedings of the 26th annual international conference on machine learning* (pp. 1105–1112). ACM.
- Wei, X., & Croft, W. B. (2006). Lda-based document models for ad-hoc retrieval. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval* (pp. 178–185). ACM.

- Wieting, J., Bansal, M., Gimpel, K., Livescu, K., & Roth, D. (2015). From paraphrase database to compositional paraphrase model and back. *Transactions of the Association for Computational Linguistics*, 3, 345–358.
- Yi, X., & Allan, J. (2009). A comparative study of utilizing topic models for information retrieval. In *Advances in Information Retrieval* (pp. 29–41). Springer Berlin Heidelberg.