



On the Coherence of Fake News Articles

Iknoor Singh¹, Deepak P.²(✉), and Anoop K.³

¹ University of Sheffield, Sheffield, UK

`iknoor.ai@gmail.com`

² Queen's University Belfast, Belfast, UK

`deepaksp@acm.org`

³ University of Calicut, Thenhpalam, India

`anoopk_dcs@uoc.ac.in`

Abstract. The generation and spread of fake news within new and online media sources is emerging as a phenomenon of high societal significance. Combating them using data-driven analytics has been attracting much recent scholarly interest. In this computational social science study, we analyze the textual coherence of fake news articles vis-a-vis legitimate ones. We develop three computational formulations of textual coherence drawing upon the state-of-the-art methods in natural language processing and data science. Two real-world datasets from widely different domains which have fake/legitimate article labellings are then analyzed with respect to textual coherence. We observe apparent differences in textual coherence across fake and legitimate news articles, with fake news articles consistently scoring lower on coherence as compared to legitimate news ones. While the relative coherence shortfall of fake news articles as compared to legitimate ones form the main observation from our study, we analyze several aspects of the differences and outline potential avenues of further inquiry.

1 Introduction

The spread of *fake news* is increasingly being recognized as a global issue of enormous significance. The phenomenon of fake news, or disinformation disguised as news, started gaining rampant attention around the 2016 US presidential elections [2]. While politics remains the domain which attracts most scholarly interest in studying the influence of fake news [7], the impact of alternative facts on economic [12] and healthcare [28] sectors are increasingly getting recognized. Of late, the news ecosystem has evolved from a small set of regulated and trusted sources to numerous online news sources and social media. Such new media sources come with limited liability for disinformation, and thus are easy vehicles for fake news. Data science methods for fake news detection within social media such as Twitter has largely focused on leveraging the social network and temporal propagation information such as response and retweet traces and their temporal build-up; the usage of the core content information within the tweet has been shallow. In fact, some recent techniques (e.g., [32]) achieve state-of-the-art performance without using any content features whatsoever. The

task landscape, however, changes significantly when one moves from the realm of tweets to online news sources (such as those in Wikipedia’s list¹); the latter form a large fraction of fake news that are debunked within popular debunking sites such as Snopes². Fake news within online news sources (e.g., [31]) are characterized by scanty network and propagation information; this is so since they are typically posted as textual articles within websites (as against originating from social media accounts), and their propagation happens on external social media websites through link sharing. Further, it is often necessary to bank on exploiting text information in order to develop fake news detection methods for scenarios such as those of a narrow scope (e.g., a local council election, or regional soccer game) even for Twitter, since the narrow scope would yield sparse and unreliable propagation information. Accordingly, there has been some recent interest in characterizing fake news in terms of various aspects of textual content. Previous work along this direction has considered satirical cues [26], expression of stance [4], rhetorical structures [25] and topical novelty [30].

In this computational social science study, we evaluate the textual coherence of fake news articles vis-a-vis legitimate ones. While our definitions of textual coherence will follow in a later section, it is quite intimately related to the notion of *cohesion* and *coherence* in language studies [21]. We choose to use the term *coherence* due to being more familiar to the computing community. Cohesion has been considered as an important feature for assessing the structure of text [21] and has been argued to play a role in writing quality [3, 6, 17].

2 Related Work

With our study being on coherence on fake news as assessed using their textual content, we provide some background into two streams of literature. First, we briefly summarize some recent literature around fake news with particular emphasis to studies directed at their lexical properties and those directed at non-lexical integrity and coherence. Second, we outline some natural language processing (NLP) techniques that we will use as technical building blocks to assess textual coherence.

Characteristics of Fake News: In the recent years, there have been abundant explorations into understanding the phenomenon of fake news. We now briefly review some selected work in the area. In probably the first in-depth effort at characterizing text content, [25] make use of rhetorical structure theory to understand differences in distribution of rhetorical relations across fake and legitimate news. They identify that disjunction and restatement appear much more commonly within legitimate stories. Along similar lines, [24] perform a stylometric analysis inquiring into whether writing style can help distinguish hyperpartisan news from others, and fake news from legitimate ones (two separate tasks).

¹ https://en.wikipedia.org/wiki/List_of_fake_news_websites.

² <http://www.snopes.com>.

While they identify significant differences in style between hyperpartisan and non-hyperpartisan articles, they observe that such differences are much more subdued across fake and legitimate news and are less likely to be useful for the latter task. In an extensive longitudinal study [30] to understand the propagation of fake and legitimate news, the authors leverage topic models [29] to quantify the divergence in character between them. They find that human users have an inherent higher propensity to spread fake news and attribute that propensity to users perceiving fake news as more novel in comparison to what they have seen in the near past, novelty assessed using topic models. Conroy et al. [5] summarize the research into identification of *deception cues* (or tell-tale characteristics of fake news) into two streams; viz., linguistic and network. While they do not make concrete observations on the relative effectiveness of the two categories of cues, they illustrate the utility of external knowledge sources such as Wikipedia in order to evaluate veracity of claims. Moving outside the realm of text analysis, there has been work on quantifying the differences between fake and legitimate news datasets in terms of image features [14]. They observe that there are often images along with microblog posts in Twitter, and these may hold cues as to the veracity of the tweet. With most tweets containing only one image each, they analyze the statistical properties of the dataset of images collected from across fake tweets, and compare/contrast them with those from a dataset of images from across legitimate tweets. They observe that the average visual coherence, i.e., the average pairwise similarities between images, is roughly the same across fake and legitimate image datasets; however, the fake image dataset has a larger dispersion of coherence scores around the mean. In devising our lexical coherence scores that we use in our study, we were inspired by the formulation of visual coherence scores as aggregates of pairwise similarities.

NLP Building Blocks for Coherence Assessments: In our study, we quantify lexical coherence by building upon advancements in three different directions in natural language processing. We briefly outline them herein.

Text Embeddings. Leveraging a dataset of text documents to learn distributional/vector representations of words, phrases and sentences has been a recent trend in natural language processing literature, due to pioneering work such as word2vec [19] and GloVe [23]. Such techniques, called word embedding algorithms, map each word in the document corpus to a vector of pre-specified dimensionality by utilizing the lexical proximity of words within the documents in which they appear. Thus, words that appear close to each other often within documents in the corpus would get mapped to vectors that are proximal to each other. These vectors have been shown to yield themselves to meaningful algebraic manipulation [20]. While word embeddings are themselves useful for many practical applications, techniques for deriving a single embedding/vector for larger text fragments such as sentences, paragraphs and whole documents [16] have been devised. That said, sentence embeddings/vectors formed by simply averaging the vectors corresponding to their component words, often dubbed average

word2vec (e.g., [8]), often are found to be competitive with more sophisticated embeddings for text fragments. In our first computational model for lexical coherence assessment, we will use average word2vec vectors to represent sentences.

Explicit Semantic Analysis. Structured knowledge sources such as Wikipedia encompass a wide variety of high-quality manually curated and continuously updated knowledge. Using them for deriving meaningful representations of text data has been a direction of extensive research. A notable technique [9] along this direction attempts to represent text as vectors in a high dimensional space formed by Wikipedia concepts. Owing to using Wikipedia concepts explicitly, it is called explicit semantic analysis. In addition to generating vector representations that is then useful for a variety of text analytics tasks, these vectors are intuitively meaningful and easy to interpret due to the dimensions mapping directly to Wikipedia articles. This operates by processing Wikipedia articles in order to derive an inverted index for words (as is common in information retrieval engines [22]); these inverted indexes are then used to convert a text into a vector in the space of Wikipedia dimensions. Explicit Semantic Analysis, or ESA as it is often referred to, has been used for a number of different applications including those relating to computing semantic relatedness [27]. In our second computation approach, we will estimate document lexical coherence using sentence-level ESA vectors.

Entity Linkings. In the realm of news articles, entities from knowledge bases such as Wikipedia often get directly referenced in the text. For example, the fragments *UK* and *European Union* in ‘*UK is due to leave the European Customs Union in 2020*’ may be seen to be intuitively referencing the respective entities, i.e., *United Kingdom* and *European Union*, in a knowledge base. Entity Linking [11] (*aka* named entity disambiguation) methodologies target to identify such references, effectively establishing a method to directly link text to a set of entities in a knowledge base; Wikipedia is the most popular knowledge base used for entity linking (e.g., [18]). Wikipedia entities may be thought of as being connected in a graph, edges constructed using hyperlinks in respective Wikipedia articles. Wikipedia graphs have been used for estimating semantic relatedness of text in previous work [34]. In short, converting a text document to a set of Wikipedia entities referenced within them (using entity linking methods) provides a semantic grounding for the text within the Wikipedia graph. We use such Wikipedia-based semantic representations in order to devise our third computational technique for measuring lexical coherence.

3 Research Objective and Computational Framework

Research Objective: Our core research objective is to study whether fake news articles differ from legitimate news articles along the dimension of *coherence as estimated using their textual content* using computational methods. We use the

term coherence to denote the overall consistency of the document in adhering to core focus theme(s) or topic(s). In a way, it may be seen as contrary to the notion of dispersion or scatter. A document that switches themes many times over may be considered as one that lacks coherence.

3.1 Computational Framework for Lexical Coherence

For the purposes of our study, we need a computational notion of coherence. Inspired by previous work within the realm of images where visual coherence is estimated using aggregate of pairwise similarities [14], we model the coherence of an article as the average/mean of pairwise similarities between ‘elements’ of the article. Depending on the computational technique, we model elements as either sentences or Wikipedia entities referenced in the article. For the sentence-level structuring, for a document D comprising d_n sentences $\{s_1, s_2, \dots, s_{d_n}\}$, coherence is outlined as:

$$C_{sent}(D) = \text{mean} \left\{ \text{sim}(\text{rep}(s_i), \text{rep}(s_j)) \mid s_i, s_j \in D, i \neq j \right\} \quad (1)$$

where $\text{rep}(s)$ denotes the representation of the sentence s , and $\text{sim}(\cdot, \cdot)$ is a suitable similarity function over the chosen representation. Two of our computational approaches use sentence level coherence assessments; they differ in the kind of representation, and consequently the similarity measure, that they use. The third measure uses entity linking to identify Wikipedia entities that are used in the text document. For a document D comprising references to d_m entities $\{e_1, e_2, \dots, e_{d_m}\}$, coherence is computed as:

$$C_{ent}(D) = \text{mean} \left\{ \text{sim}(e_i, e_j) \mid e_i, e_j \in D, i \neq j \right\} \quad (2)$$

where $\text{sim}(\cdot, \cdot)$ is a suitable similarity measure between pairs of Wikipedia entities. We consistently use numeric vectors as representations of sentences and entities, and employ cosine similarity³, the popular vector similarity measure, to compute similarities between vectors.

3.2 Profiling Fake and Legitimate News Using Lexical Coherence

Having outlined our framework for computing document-specific lexical coherence, we now describe how we use it in order to understand differences between fake and legitimate news articles. Let $\mathcal{F} = \{\dots, F, \dots\}$ and $\mathcal{L} = \{\dots, L, \dots\}$ be separate datasets of fake news and legitimate news articles respectively. Each document in \mathcal{F} and \mathcal{L} is subjected to the coherence assessment, yielding a single document-specific score (for each computational approach for quantifying lexical coherence). These yield separate sets of lexical coherence values for \mathcal{F} and \mathcal{L} :

$$C(\mathcal{F}) = \{\dots, C(F), \dots\} \quad C(\mathcal{L}) = \{\dots, C(L), \dots\} \quad (3)$$

³ https://en.wikipedia.org/wiki/Cosine_similarity.

where $C(\cdot)$ is either of C_{sent} or C_{ent} . Aggregate statistics across the sets $C(\mathcal{F})$ and $C(\mathcal{L})$, such as mean and standard deviation, enable quantifying the relative differences in textual coherence between fake and legitimate news.

4 Computational Approaches

We now outline our three computational approaches to quantify coherence at the document level. The first and second approaches use sentence-level modelling, and leverage word embeddings and explicit semantic analysis respectively. The third uses entity linking to convert each document into a set of entities, followed by computing the coherence at the entity set level. Each coherence quantification method is fully specified with the specification of how the vector representation is derived for the elements (entities or sentences) of the document.

Coherence Using Word Embeddings. This is the first of our sentence-level coherence assessment methods. As outlined in the related work section, word2vec [19] is among the most popular word embedding methods. Word2vec uses a shallow, two-layer neural network that is trained to reconstruct linguistic contexts of words over a corpus of documents. Word2vec takes as its input a large corpus of text and produces a vector space, typically of several hundred dimensions, with each unique word in the corpus being assigned a corresponding vector in the space. We use a pre-trained word2vec vector dataset that was trained over a huge corpus⁴ since they are likely to be better representations being learnt over a massive dataset. Each sentence s_i in document D is then represented as the average of the word2vec vectors of the words, denoted $\text{word2vec}(w)$, it contains:

$$\text{rep}(s_i) = \text{mean}\{\text{word2vec}(w) | w \in s_i\} \quad (4)$$

This completes the specification of coherence quantification using embeddings.

Coherence Using Explicit Semantic Analysis. Explicit Semantic Analysis (ESA) [9] forms the basis of our second sentence-level coherence quantification method. ESA starts with a collection of text articles sources from a knowledge base, typically Wikipedia; each article is turned into a *bag of words*. Each word may then be thought of as being represented as a vector over the set of Wikipedia articles, each element of the vector directly related to the number of times it appears in the respective article. The ESA representation for each sentence is then simply the average of the ESA vectors of the words that it contains.

$$\text{rep}(s_i) = \text{mean}\{\text{esa}(w) | w \in s_i\} \quad (5)$$

where $\text{esa}(w)$ is the vector representation of the word w under ESA.

⁴ <https://github.com/mmihaltz/word2vec-GoogleNews-vectors>.

Coherence Using Entity Linking. Given a document, entity linking (EL) methods [10] identify mentions of entities within them. Identifying Wikipedia entities to associate references to, is often complex and depends on the context of the word. EL algorithms use a variety of different heuristics in order to accurately identify the set of entities referenced in the document. We would like to now convert each entity thus identified to a vector so that it may be used for document coherence quantification within our framework. Towards this, we use the Wikipedia2vec technique [33] which is inspired by Word2vec and forms vectors for Wikipedia entities by processing the corpus of Wikipedia articles. Thus, the representation of each entity is simply the Wikipedia2vec vector associated with that entity. That representation then feeds into Eq. 2 for coherence assessments.

5 Experimental Study

5.1 Datasets

In order to ensure the generalizability of the insights from our coherence study, we evaluate the coherence scores over two datasets. The first one, *ISOT fake news dataset*, is a publicly available dataset comprising 10k+ articles focused on politics. The second dataset is one sourced by ourselves comprising 1k articles on *health and well-being* (HWB) from various online sources. These datasets are very different both in terms of size and the kind of topics they deal with. It may be noted that the nature of our task makes datasets comprising long text articles more suitable. Most fake news datasets involve tweets, and are thus sparse with respect to text data⁵, making them unsuitable for textual coherence studies such as ours. Thus, we limit our attention to the aforementioned two datasets both of which comprise textual articles. We describe them separately herein.

ISOT Fake News Dataset. The ISOT Fake News dataset⁶ [1] is the largest public dataset of textual articles with fake and legitimate labellings that we have come across. The ISOT dataset comprises various categories of articles, of which the *politics* category is the only one that appears within both *fake* and *real/legitimate* labellings. Thus, we use the politics subset from both fake and legitimate categories for our study. The details of the dataset as well as the statistics from the sentence segmentation and entity linking appear in Table 1.

HWB Dataset. We believe that health and well-being is another domain that is often targeted by fake news sources. As an example, fake news on topics such as vaccinations has raised significant concerns [13] in recent times, not to mention the COVID-19 pandemic. Thus, we curated a set of articles with fake/legitimate labellings tailored to the health domain. For the legitimate news articles, we crawled 500 news documents on health and well-being from reputable sources

⁵ Tweets are limited to a maximum of 280 characters.

⁶ <https://www.uvic.ca/engineering/ece/isot/datasets/index.php>.

Table 1. Dataset Statistics (SD = Standard Deviation)

Dataset	Category	#Articles	#Sentences Per Article Mean (SD)	#Entities Per Article Mean (SD)
ISOT	Fake	5816	11.56 (10.54)	38.99 (30.17)
	Legitimate	11175	14.65 (10.71)	46.48 (27.31)
HWB	Fake	500	24.62 (18.04)	49.56 (28.67)
	Legitimate	500	28.24 (19.37)	62.91 (37.07)

such as CNN, NYTimes, Washington Post and New Indian Express. For fake news, we crawled 500 articles on similar topics from well-reported misinformation websites such as BeforeItsNews, Nephew and MadWorldNews. These were manually verified for category suitability, thus avoiding blind reliance on source level labellings. This dataset, which we will refer to as HWB, short for *health and well-being*, will be made available at <https://dcs.uoc.ac.in/cida/resources/hwb.html>. HWB dataset statistics also appear in Table 1.

On the Article Size Disparity. The disparity in article sizes between fake and legitimate news articles is important to reflect upon, in the context of our comparative study. In particular, it is important to note how coherence assessments may be influenced by the number of sentences and entity references within each article. It may be intuitively expected that coherence quantification would yield a lower value for longer documents than for shorter ones, given that all pairs of sentences/entities are used in the comparison. Our study stems from the hypothesis that fake articles may be *less* coherent; the article length disparity in the dataset suggests that the null hypothesis assumption (that coherence is similar across fake and legitimate news) being true would yield *higher* coherence score for the fake news documents (being shorter). *In a way, it may be observed that any empirical evidence illustrating lower coherence scores for fake articles, as we observe in our results, could be held to infer a stronger departure from the null hypothesis than in a dataset where fake and legitimate articles were of similar sizes.* Thus, the article length distribution trends only deepen the significance of our results that points to lower coherence among fake articles.

5.2 Experimental Setup

We now describe some details of the experimental setup we employed. The code was written in Python. NLTK⁷, a popular natural language toolkit, was used for sentence splitting and further processing. The word embedding coherence assessments were performed using Google’s pre-trained word2vec corpus⁸, which was trained over news articles. Explicit Semantic Analysis (ESA) was run using the publicly available EasyESA⁹ implementation. For the entity linking method,

⁷ <https://www.nltk.org/>.

⁸ <https://code.google.com/archive/p/word2vec/>.

⁹ <https://github.com/dscarvalho/easyesa>.

Table 2. Coherence Results Summary (SD = Standard Deviation). Statistically significant results with $p < 0.05$ (two-tailed t-test) in bold. XE-Y is a commonly used mathematical abbreviation to stand for $X \times 10^{-Y}$

Dataset	Category	Word Embedding Coherence Mean (SD)	ESA Coherence Mean (SD)	Entity Linking Coherence Mean (SD)
ISOT	Fake	0.546518 (0.0070)	0.999218 (4.90E-05)	0.277454 (0.0010)
	Legitimate	0.567870 (0.0055)	0.999474 (2.45E-05)	0.286689 (0.0003)
	Difference in %	3.91%	0.03%	3.33%
	p-value	6.29E-60	0.013341	1.90E-100
HWB	Fake	0.468907 (0.0055)	0.995245 (0.0040)	0.307874 (0.0006)
	Legitimate	0.506322 (0.0059)	0.997276 (0.0020)	0.318574 (0.0008)
	Difference in %	7.98%	0.20%	3.48%
	p-value	1.46E-14	0.557432	6.27E-10

the named entities were identified using the NLTK toolkit, and their vectors were looked up on the Wikipedia2vec¹⁰ pre-trained model¹¹.

5.3 Analysis of Text Coherence

We first analyze the mean and standard deviation of coherence scores of fake and legitimate news articles as assessed using each of our three methods. Higher coherence scores indicate higher textual coherence. Table 2 summarizes the results. *In each of the three coherence assessments across two datasets, thus six combinations overall, the fake news articles were found to be less coherent than the legitimate news ones on the average.* The difference was found to be statistically significant with $p < 0.05$ under the two-tailed t-test¹² in five of six combinations, recording very low p-values (i.e., strong difference) in many cases.

Trends Across Methods. The largest difference in means are observed for the word embedding coherence scores, with the legitimate news articles being around 4% and 8% more coherent than fake news articles in the ISOT and HWB datasets respectively. The lowest p-values are observed for the entity linking method, where the legitimate articles are 3+% more coherent than fake news articles across datasets for the ISOT dataset; the p-value being in the region of 1E-100 indicates the presence of consistent difference in coherence across fake and legitimate news articles. On the other hand, the coherence scores for ESA vary only slightly in magnitude across fake and legitimate news articles. This is likely because ESA is primarily intended to separate articles from different domains; articles within the same domain thus often get judged to be very close to each other, as we will see in a more detailed analysis in the later section. The smaller differences across fake and legitimate news articles are still statistically significant for the ISOT dataset (p-value < 0.05), whereas they are not so in the

¹⁰ <https://github.com/wikipedia2vec/wikipedia2vec>.

¹¹ <https://wikipedia2vec.github.io/wikipedia2vec/pretrained/>.

¹² <https://www.itl.nist.gov/div898/handbook/eda/section3/eda353.htm>.

case of the much smaller HWB dataset. It may be appreciated that statistical significance assessments depend on degrees of freedom, roughly interpreted as the number of independent samples; this makes it harder to approach statistical significance in small datasets. Overall, these results also indicate that the word embedding perspective is best suited, among the three methods, to discern textual coherence differences between legitimate and fake news.

Coherence Spread. The spread of the coherence across articles was seen to be largely broader for fake news articles, as compared to legitimate ones. This is evident from the higher standard deviations exhibited by the coherence scores in the majority of cases. From observations over the dataset, we find that the coherence scores for legitimate articles generally form a unimodal distribution, whereas the distribution of coherence scores across fake news articles show some minor deviations from unimodality. In particular, we find a small number of scores clustered in the low range of the spectrum, and another much larger set of scores forming a unimodal distribution centered at a score lesser than that the respective centre for the legitimate news articles. This difference in character, of which more details follow in the next section, reflects in the higher standard deviation in coherence scores for the fake news documents.

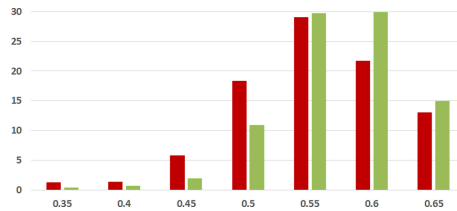


Fig. 1. ISOT Word Embedding Coherence Histogram (%articles vs. coherence scores): Legitimate and Fake articles’ scores in green and red respectively. (Color figure online)

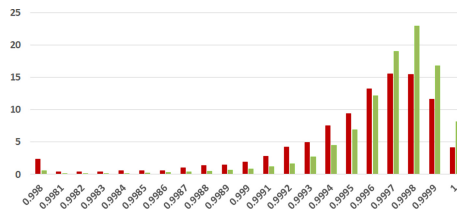


Fig. 2. ISOT ESA Coherence Histogram (%articles vs. coherence scores): Legitimate and Fake articles’ scores in green and red respectively. (Color figure online)

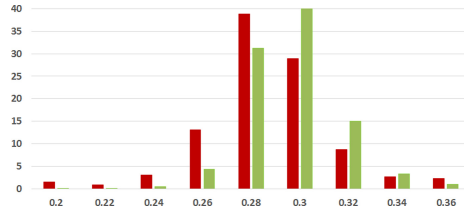


Fig. 3. ISOT Entity Linking Coherence Histogram (%articles vs. coherence scores): Legitimate and Fake articles' scores in green and red respectively. (Color figure online)

Coherence Score Histograms. We visualize the nature of the coherence score distributions for the various methods further by way of histogram plots in Figs. 1, 2 and 3 for the ISOT dataset. The corresponding histograms for the HWB dataset appear in Figs. 4, 5 and 6. We have set the histogram buckets in a way to amplify the region where most documents fall, density of scores being around different ranges for different methods. This happens to be around 0.5 for word embedding scores, 0.999 for ESA scores, and around 0.3 for entity linking scores. With the number of articles differing across the fake and legitimate subsets, we have set the Y-axis to indicate the percentage of articles in each of the ranges (as opposed to raw frequency counts), to aid meaningful visual comparison. All documents that fall outside the range in the histogram are incorporated into the leftmost or rightmost pillar in the histogram as appropriate.

The most striking high-level trend, across the six histograms is as follows. When one follows the histogram pillars from left to right, the red pillar (corresponding to fake news articles) is consistently taller than the green pillar (corresponding to legitimate news articles), until a point beyond which the trend reverses; from that point onwards the green pillar is consistently taller than the red pillar. Thus, the lower coherence scores have a larger fraction of fake articles than legitimate ones and vice versa. For example, this point of reversal is at 0.55 for Fig. 1, 0.9997 for Fig. 3 and 0.9996 for Fig. 5. There are only two histogram points that are not very aligned with this trend, both for the entity linking method, which are 0.36 in Fig. 3 and 0.34 in Fig. 6; even in those cases, the high-level trend is still consistent with our analysis.

The second observation, albeit unsurprising, is the largely single-peak (i.e., unimodal) distribution of the coherence score in each of the six charts for both fake and legitimate news articles; this also vindicates our choice of the statistical significance test statistic in the previous section, t-test being suited best for comparing unimodal distributions. A slight departure from that unimodality, as alluded to earlier, is visible for fake article coherence score. This is most expressed for the ESA method, with a the leftmost red pillar being quite tall in Figs. 2 and 5; it may be noted that the leftmost pillars count all documents below that score and thus, the small peak that exists in the lower end of the fake news scores is overemphasized in the graphs due to the nature of the plots.

Thirdly, the red-green reversal trend as observed earlier, may be interpreted as largely being an artifact of the relative positioning of the centres of the unimodal score distribution across fake and legitimate news articles. The red peak appears to the left (i.e., at a lower coherence score) of the green peak; it is easy to observe this trend if one looks at the red and green distributions separately on the various charts. For example, the red peak for Fig. 1 is at 0.55, whereas the green peak is at 0.60. Similarly, the red peak in Fig. 3 is at 0.28, whereas the green peak is at 0.30. Similar trends are easier to observe for the HWB results.

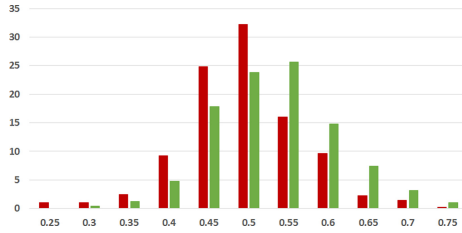


Fig. 4. HWB Word Embedding Coherence Histogram (%articles vs. coherence scores): Legitimate and Fake articles’ scores in green and red respectively. (Color figure online)

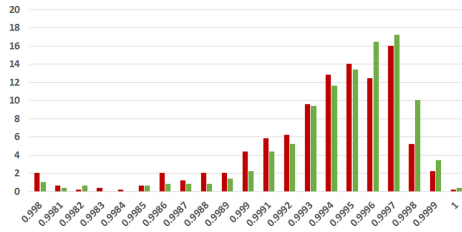


Fig. 5. HWB ESA Coherence Histogram (%articles vs. coherence scores): Legitimate and Fake articles’ scores in green and red respectively. (Color figure online)

Summary. Main observations from our computational social science study are:

- *Fake news articles less coherent:* The trends across all the coherence scoring mechanisms over the two widely different datasets (different domains, different sizes) indicate that fake news articles are less coherent than legitimate news ones. The trends are statistically significant in all but one case.
- *Word embeddings most suited:* From across our results, we observe that word embedding based mechanism is most suited, among our methods, to help discern the difference in coherence between fake and legitimate news articles.

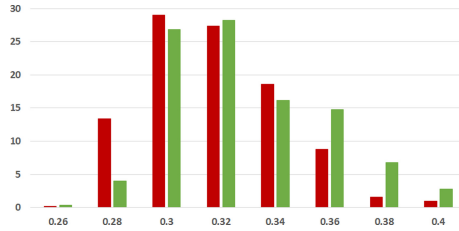


Fig. 6. HWB Entity Linking Coherence Histogram (%articles vs. coherence scores): Legitimate and Fake articles' scores in green and red respectively. (Color figure online)

- *Unimodal distributions with different peaks*: The high-level trend points to a unimodal distribution of coherence scores (slight departure observed for fake articles), with the score distribution for fake news peaking at a lower score.

6 Discussion

The relevance or importance of this computational social science study lies in what the observations and the insights from them point us to. We discuss some such aspects in this section.

6.1 Towards More Generic Fake News Detection

Fake news detection has been traditionally viewed as a classification problem by the data science community, with the early approaches relying on the availability of labelled training data. The text features employed were standard ones used within the text mining/NLP community, such as word-level ones and coarser lexical features. Standard text classification scenarios such as identifying the label to be associated with a scholarly article or determining disease severity from medical reports involve settings where the human actors involved in generating the text are largely passive to the classification task. On the other hand, data science based fake news identification methods stand between the fake news author and accomplishment of his/her objectives (which may be one of myriad possibilities such as influencing the reader's political or medical choices). Social networks regularly use fake news detection algorithms to aid prioritization of stories. Making use of low-level features such as words and lexical patterns in fake news detection makes it easier for the fake news author to circumvent the fake news filter and reach a broad audience. As an example, a fake news filter that is trained on US based political corpora using word-level features could be easily surpassed by using words and phrases that are not commonly used within US politics reporting (e.g., replacing *President* with *head of state*). On the other hand, moving from low-level lexical features (e.g., words) to higher level ones such as topical novelty (as investigated in a Science journal article [30]), emotions [15]

and rhetorical structures [25] would yield more ‘*generic*’ fake news identification methods that are more robust to being easily tricked by fake news peddlers. We observe that moving to higher-level features for fake news identification has not yet been a widespread trend within the data science community; this is likely to be due to the fact that investigating specific trends do not necessarily individually improve the state-of-the-art for fake news detection using conventional metrics such as empirical accuracy that are employed within data science. Nonetheless, such efforts yield insights which hold much promise in collectively leading to more generic fake news detection for the future. Techniques that use high-level features may also be better transferable across domains and geographies. We believe that our work investigates an important high-level feature, that of lexical coherence, and provides insights highly supported by datasets across widely varying domains, and would be valuable in contributing to a general fake news detection method less reliant on abundant training datasets.

6.2 Lexical Coherence and Other Disciplines

We now consider how our observations could lead to interesting research questions in other disciplinary domains.

News Media. The observation that fake news articles exhibit lower textual coherence could be translated into various questions when it comes to news media. It may be expected that fake news articles appear in media sources that are not as established, big or reputed as the ones that report only accurate news; these media houses may also have less experienced staff. If such an assumption is true, one could possibly consider various media-specific reasons for the relative coherence trends between fake and legitimate news:

- Is the lower coherence of fake news articles a reflection of less mature media quality control employed within smaller and less established media houses?
- Is the relative lack of coherence of fake news articles more reflective of amateurish authors with limited journalistic experience?

Linguistics. Yet another perspective to analyze our results is that of linguistics, or more specifically, cognitive linguistics¹³. The intent of fake news authors are likely more towards swaying the reader towards a particular perspective/stance/position; this is markedly different from just conveying accurate and informative news, as may be expected of more reputed media. Under the lens of cognitive linguistics, it would be interesting to analyze the following questions:

- Is the reduced coherence due to fake news generally delving into multiple topics within the same article?
- Could the reduced coherence be correlated with mixing of emotional and factual narratives?

¹³ https://en.wikipedia.org/wiki/Cognitive_linguistics.

Besides the above, our results may spawn other questions based on perspectives that we are unable to view it from, given the limited breadth in our expertise.

7 Conclusions and Future Work

We studied the variation between fake and legitimate news articles in terms of the coherence of their textual content. Within the computational social science framework, we used state-of-the-art data science methodologies to formulate three different computational notions of textual coherence, in order to address our research question. The methods widely vary in character and use word embeddings, explicit semantic analysis and entity linking respectively. We empirically analyzed two datasets, one public dataset from politics, and another one comprising health and well-being articles, in terms of their textual coherence, and analyzed differences across their fake and legitimate news subsets. All the results across all six combinations (3 scoring methods, 2 datasets) unequivocally indicated that fake news article subset exhibits lower textual coherence as compared to the legitimate news subset. These results are despite the fact that fake news articles were found to be shorter, short articles intuitively having a higher propensity to be more coherent. In summary, our results suggest that fake news articles are less coherent than legitimate news articles in systematic and discernible ways when analyzed using simple textual coherence scoring methods such as the one we have devised.

Acknowledgements. Deepak P was supported by MHRD SPARC (P620).

References

1. Ahmed, H., Traore, I., Saad, S.: Detection of online fake news using N-gram analysis and machine learning techniques. In: Traore, I., Woungang, I., Awad, A. (eds.) ISDDC 2017. LNCS, vol. 10618, pp. 127–138. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-69155-8_9
2. Allcott, H., Gentzkow, M.: Social media and fake news in the 2016 election. *J. Econ. Perspect.* **31**, 211–36 (2017)
3. Alotaibi, H.: The role of lexical cohesion in writing quality. *Int. J. Appl. Linguist. Engl. Lit.* **4**(1), 261–269 (2015)
4. Chopra, S., Jain, S., Sholar, J.M.: Towards automatic identification of fake news: headline-article stance detection with LSTM attention models (2017)
5. Conroy, N.J., Rubin, V.L., Chen, Y.: Automatic deception detection: methods for finding fake news. In: Proceedings of the 78th ASIS&T Annual Meeting, p. 82 (2015)
6. Crossley, S., McNamara, D.: Text coherence and judgments of essay quality: models of quality and coherence. In: Proceedings of the Annual Meeting of the Cognitive Science Society, vol. 33 (2011)
7. Davies, W.: The age of post-truth politics. *The New York Times* **24**, 2016 (2016)
8. Dilawar, N., et al.: Understanding citizen issues through reviews: a step towards data informed planning in smart cities. *Appl. Sci.* **8**, 1589 (2018)

9. Gabrilovich, E., Markovitch, S., et al.: Computing semantic relatedness using Wikipedia-based explicit semantic analysis. In: *IJCAI*, vol. 7, pp. 1606–1611 (2007)
10. Gupta, N., Singh, S., Roth, D.: Entity linking via joint encoding of types, descriptions, and context. In: *Proceedings of the Conference on EMNLP*, pp. 2681–2690 (2017)
11. Hoffart, J., et al.: Robust disambiguation of named entities in text. In: *Proceedings of the Conference on EMNLP*, pp. 782–792 (2011)
12. Hopkin, J., Rosamond, B.: Post-truth politics, bullshit and badideas: ‘deficit fetishism’ in the UK. *New Polit. Econ.* **23**, 641–655 (2018)
13. Iacobucci, G.: Vaccination: “fake news” on social media may be harming UK uptake, report warns. *BMJ: Br. Med. J. (Online)* **364** (2019)
14. Jin, Z., Cao, J., Zhang, Y., Zhou, J., Tian, Q.: Novel visual and statistical image features for microblogs news verification. *IEEE Trans. Multimed.* **19**, 598–608 (2017)
15. Anoop, K., Deepak, P., Lajish, V.L.: Emotion cognizance improves health fake news detection. In: *24th International Database Engineering & Applications Symposium (IDEAS 2020)* (2020)
16. Le, Q., Mikolov, T.: Distributed representations of sentences and documents. In: *International Conference on Machine Learning*, pp. 1188–1196 (2014)
17. McCulley, G.A.: Writing quality, coherence, and cohesion. *Res. Teach. Engl.* 269–282 (1985)
18. Mihalcea, R., Csosmai, A.: Wikify!: linking documents to encyclopedic knowledge. In: *CIKM*, pp. 233–242 (2007)
19. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: *NIPS* (2013)
20. Mikolov, T., Yih, W., Zweig, G.: Linguistic regularities in continuous space word representations. In: *NAACL*, pp. 746–751 (2013)
21. Morris, J., Hirst, G.: Lexical cohesion computed by thesaural relations as an indicator of the structure of text. *Comput. Linguist.* **17**(1), 21–48 (1991)
22. Ounis, I., Amati, G., Plachouras, V., He, B., Macdonald, C., Johnson, D.: Terrier information retrieval platform. In: Losada, D.E., Fernández-Luna, J.M. (eds.) *ECIR 2005. LNCS*, vol. 3408, pp. 517–519. Springer, Heidelberg (2005). https://doi.org/10.1007/978-3-540-31865-1_37
23. Pennington, J., Socher, R., Manning, C.: Glove: global vectors for word representation. In: *Proceedings of the Conference on EMLNP*, pp. 1532–1543 (2014)
24. Potthast, M., Kiesel, J., Reinartz, K., Bevendorff, J., Stein, B.: A stylometric inquiry into hyperpartisan and fake news. [arXiv:1702.05638](https://arxiv.org/abs/1702.05638) (2017)
25. Rubin, V.L., Conroy, N.J., Chen, Y.: Towards news verification: deception detection methods for news discourse. In: *Hawaii International Conference on System Sciences* (2015)
26. Rubin, V., Conroy, N., Chen, Y., Cornwell, S.: Fake news or truth? Using satirical cues to detect potentially misleading news. In: *Proceedings of the 2nd Workshop on Computational Approaches to Deception Detection*, pp. 7–17 (2016)
27. Scholl, P., Böhnstedt, D., Domínguez García, R., Rensing, C., Steinmetz, R.: Extended explicit semantic analysis for calculating semantic relatedness of web resources. In: Wolpers, M., Kirschner, P.A., Scheffel, M., Lindstaedt, S., Dimitrova, V. (eds.) *EC-TEL 2010. LNCS*, vol. 6383, pp. 324–339. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-16020-2_22
28. Speed, E., Mannon, R.: The rise of post-truth populism in pluralist liberal democracies: challenges for health policy. *Int. J. Health Policy Manag.* **6**, 249 (2017)
29. Steyvers, M., Griffiths, T.: Probabilistic topic models. *Handb. Latent Semant. Anal.* **427**, 424–440 (2007)

30. Vosoughi, S., Roy, D., Aral, S.: The spread of true and false news online. *Science* **359**, 1146–1151 (2018)
31. Wargadiredja, A.T.: Indonesian teens are getting ‘drunk’ off boiled bloody menstrual pads (2018). https://www.vice.com/en_asia/article/wj38gx/indonesian-teens-drinking-boiled-bloody-menstrual-pads. Accessed Apr 2020
32. Wu, L., Liu, H.: Tracing fake-news footprints: characterizing social media messages by how they propagate. In: *WSDM*, pp. 637–645 (2018)
33. Yamada, I., Asai, A., Shindo, H., Takeda, H., Takefuji, Y.: Wikipedia2vec: an optimized implementation for learning embeddings from wikipedia. [arXiv:1812.06280](https://arxiv.org/abs/1812.06280) (2018)
34. Yeh, E., Ramage, D., Manning, C.D., Agirre, E., Soroa, A.: Wikiwalk: random walks on Wikipedia for semantic relatedness. In: *Proceedings Workshop on Graph-based Methods for NLP*, pp. 41–49. *ACL* (2009)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

