

Distinguishing Commercial from Editorial Content in News

Timo Kats¹[0000-0003-1650-1814], Peter van der Putten¹[0000-0002-6507-6896],
and Jasper Schelling²[0000-0002-9995-1505]

¹ Leiden University, Niels Bohrweg 1, 2333 CA Leiden, The Netherlands
t.p.a.kats@liacs.leidenuniv.nl, p.w.h.van.der.putten@liacs.leidenuniv.nl

² Stichting ACED, Amsterdam, The Netherlands
jasper@aced.site

Abstract. How can we distinguish commercial from editorial content in news, or more specifically, differentiate between advertorials and regular news articles? An advertorial is a commercial message written and formatted as an article, making it harder for readers to recognize these as advertising, despite the use of disclaimers. In our research we aim to differentiate the two using a machine learning model, and a lexicon derived from it. This was accomplished by scraping 1.000 articles and 1.000 advertorials from four different Dutch news sources and classifying these based on textual features. With this setup our most successful machine learning model had an accuracy of just over 90%. To generate additional insights into differences between news and advertorial language, we also analyzed model coefficients and explored the corpus through co-occurrence networks and t-SNE graphs.

Keywords: advertorials · NLP · t-SNE · co-occurrence networks

1 Introduction

In journalism it is best practice to clearly distinguish between editorial and sponsored commercial content. This is referred to as the ‘separation of church and state’ in media [2]. However, some forms of advertising have made this separation less clear to readers and therefore threaten this principle.

An example of this is the advertorial, which is commercial content in the form of an article. Advertorials are an example of what marketers call ‘native advertising’. In fact, advertorials are so much like articles, that despite using disclaimers and different layouts most readers don’t notice the difference. In a study conducted by the university of Georgia only 8% of readers recognized advertorials as commercial content [16]. As a result of this, advertorials have made the separation of church and state in the news less clear.

That’s why this research aims to differentiate articles and advertorials using machine learning. We would like to answer two research questions. Firstly, to what extent can we differentiate commercial and editorial content by a using machine learning model, and a lexicon derived from this? Secondly, can we use

AI and machine learning to better understand the difference between commercial and editorial language?

It's important to note that the separation of commercial and editorial content is hotly debated in journalism and the society at large. Yet, to our knowledge a machine learning based perspective to identify advertorials and commercial messaging was not part of this debate yet. By doing this we not only hope to answer our research questions, but also showcase how machine learning can be a solution in the debate surrounding the usage of advertorials. This research has been carried out in the context of the Reverb Channel program [12], a data driven exploration of our networked news culture that aims to reverse the sometimes questionable role of AI in digital media, by using it to investigate topics such as framing, polarization and ideology spaces.³

The remainder of this paper is structured as follows. Section 2 provides more background and related work. Section 3 explains the process of acquiring our data, followed by sections on our classification approach, and on our exploratory co-occurrence network based approach to increase the insight into how language differs across advertorials and news. Section 6 concludes the paper.

2 Background and Related Work

Even though we are not aware of any other research to leverage machine learning to distinguish advertorials from editorial content, the discussion around the usage of advertorials and commercial content in general is broader than this research alone, and has been debated widely in journalism and marketing. In this section we discuss some of this background context.

2.1 The change of journalism's business model in the digital age.

The rise of the internet has had a lot of effect on journalism. It opened up a whole new channel for news content, but it also negatively impacted circulation and advertising revenue for traditional news channels. For example, US weekend circulation of newspapers declined from 59.4 million (2000) to est. 25.8 million (2020), revenue from advertising declined from 48.7 billion (2000) to est. 8.8 billion (2020), whilst revenue from circulation remained relatively stable (10.5 (2000) to 11.1 (2020)), and the share of advertising revenue increased from 17% (2011) to 39% (2020) [10]. So despite drastic drops in circulation, companies were able to protect circulation income, but advertising revenues dropped dramatically. These developments altered the business model of journalism significantly, and drove publishers to find new sources of advertising revenue, such increased usage of advertorials and other forms of sponsored content.

2.2 Disguise, deception and disclosure in advertorials.

As discussed, whilst in journalism the distinction between editorial and sponsored commercial content is a key principle, this is challenged by advertorials in

³ <https://www.aced.site/en/programmes/reverb-channel>

practice as readers have a hard time differentiating these from editorial content, despite the use of labels and disclaimers [2].

Advertorials can be both deceptive and effective. As a classical example, in 1989 the R.J Reynolds Tobacco Company had settled charges with the FTC on that it had made false and misleading claims in an advertorial on health effects of smoking, titled ‘Of cigarettes and science’. Wilkinson et al. subsequently ran a test and over a quarter of participants thought the article was editorial content, not commercial [15].

In another study by Kim et al., the use of an advertorial over a standard advertisement increased the relevance of and attention to the message, and message and elaboration and recall. It made no difference whether the advertorials were labeled as such, and over two thirds of subjects exposed to labeled advertorials were not able to recall whether these advertorials were labeled or not [7].

As mentioned in the introduction, in another study by the University of Georgia only 8% of readers recognized advertorials as commercial content [16]. In their study, the use of disclaimers did have a positive impact on recognizing the text as commercial, with best effects for placement of disclaimers in the middle or the bottom, and explicit use of words such as ‘advertizing’ and ‘sponsored’. Also Krouwer et al. found that small changes, such as the location of a disclaimer, significantly impacts the recognizability for readers [8]. Apart from readers not noticing labeling, advertorials often violate guidelines for labeling, formatting and content [1].

To provide perhaps a somewhat more positive view on advertorials, in a survey by Reijmersdal et al. of subscribers of Dutch magazines, when asked explicitly only 12% of respondents thought advertorials are deceptive [11].

The more established newspapers and magazines will make more of a serious effort to make it known that certain content is sponsored, and writers producing advertorials are kept separate from the editorial teams. But is that sufficient, also when taking the proliferation of new digital media titles and the ongoing pressure to increase advertizing revenues into account, and norms are shifting towards further integration between editorial and commercial teams and objectives [3]?

The results above may vary but in our opinion this is clearly not sufficient. The ability to disguise content, willingly or unwillingly, and the probability that advertorials are not recognized as such even if properly labelled is significant. Marketers call it native advertizing for a reason.

The risk of mistaking commercial content for objective editorial content is somewhat obvious, but note there can be an opposite detrimental effect as well. For instance, Iversen at al. observed that exposure to native political ads reduced the public’s trust in political news [5].

2.3 The usage of lexicons in classifying text

As mentioned, we aim to create a classification model and lexicon that distinguishes editorial from commercial language. Whilst text classification models are used abundantly in NLP research, we are also looking to distribute our artifacts

to journalists and other non-technical audiences. In domains such as social science lexicons are often used, for common tasks such as sentiment analysis [13] or more specific tasks, such as detecting moral foundations in ethical reasoning [4, 14]. Lexicons can be handcrafted or created through linguistic analysis, and typically include keywords that indicate a particular class, potentially including a weight.

We were not able to identify prior work that uses machine learning, handcrafted or trained lexicons to differentiate advertorials from editorial content. A different, yet relevant related work is the study by Zhou, who uses genre analysis to characterize the general structure and linguistic characteristics of advertorials, using mostly manual analysis and interpretation [17].

3 Data Acquisition

In order to make a model that answers the research questions mentioned earlier we have created a data set with advertorials and regular news articles. The Reverb Channel corpus contains millions of articles [12], but no advertorials, hence we had to acquire our own data for this research. In this section we explain this process and showcase the data set that we acquired. For full details we refer to [6].

3.1 Scraping the data

The data for this research had to be scraped directly from news sources using web crawlers. For our research we used Python and the BeautifulSoup library. With this set up we made a URL-scrapers and a web-scrapers for every news source. We first collected the URLs from the pages we wanted to scrape data from and thereafter use those URLs to collect all the data we needed with the web-scrapers. We also carried out additional cleaning and transformation, such as removal of all commas, translation of any HTML to flat text where needed, and lowercasing of all text.⁴

3.2 Resulting data set

The data set that we acquired with this method has 2000 entries in total, about half of these entries are advertorials (see Figure 1). These entries are roughly equally distributed over four different news sources. These news sources are (online-only news) Nu.nl, (politically conservative) Telegraaf, (politically progressive) NRC, and (business publication) De Ondernemer. By including these four different news sources with roughly equal number of documents in the data set we strive to create an unbiased data set that is representative of the Dutch media landscape as a whole.

⁴ Source code, the lexicon and other deliverables can be found at https://github.com/TimoKats/research_distinguishing_commercial_and_editorial_content




Fig. 1: Metadata from the acquired data set

4 Distinguishing Advertorials from Regular Articles with Classification Models

With this corpus we developed classification models and a corresponding lexicon, which also gave us some first insights into differences between the language used in advertorials versus news.

4.1 Experimental set up

In terms of cleaning the data, we first removed potential leakers. Leaking variables in our model refer to words that trigger the model whilst being unique to our data set and media covered, for example sponsor names and disclaimers. To further lower the risk of leakage, we excluded the title and focused on the main text. Furthermore, we experimented with regular bag of words (BoW) as well as TFIDF weighted BoW, the removal of stop words and the number of features. Obviously, we could have easily obtained classifiers with near perfect performance, for instance by including disclaimer texts, but we were primarily interested in models that could distinguish commercial from editorial language.

For modeling, we selected a diverse set of classification methods to experiment with: SVM (default with rbf kernel), linearSVC, decision tree, random forest, k-NN, SGD and naive bayes. We restricted ourselves to these more classical methods as opposed to deep learning methods such as BERT, given that our data sets are relatively small, and interpretability of the results is key, for instance to iteratively identify leakers and get more insights into the difference between text types.

To simplify the approach, we aim to find the best performing model (incl. parameter optimization) through narrowing down the search as the experiment progresses, taking the best performing preliminary results and continuing to optimize it. A limitation of such an approach is that the estimate of final accuracy may be somewhat optimistic given the sequential nature of the experiments (manual overfitting), but a full multidimensional experimental set up was too

representation	learning model	accuracy	f1 score	auc
bag of words	svm	0.85±0.04	0.85±0.05	0.93±0.04
bag of words	linearSVC	0.84±0.05	0.84±0.06	0.9±0.05
bag of words	decisionTree	0.78±0.08	0.78±0.08	0.79±0.07
bag of words	randomForest	0.88±0.06	0.89±0.07	0.94±0.05
bag of words	k-NN	0.57±0.14	0.63±0.12	0.58±0.16
bag of words	SGD	0.87±0.07	0.86±0.08	0.93±0.05
bag of words	naiveBayes	0.76±0.11	0.77±0.09	0.76±0.11
tfidf	svm	0.89±0.05	0.89±0.05	0.94±0.04
tfidf	linearSVC	0.91±0.05	0.91±0.05	0.95±0.03
tfidf	decisionTree	0.78±0.07	0.79±0.07	0.8±0.07
tfidf	randomForest	0.88±0.07	0.89±0.06	0.94±0.05
tfidf	k-NN	0.51±0.03	0.64±0.02	0.51±0.05
tfidf	SGD	0.9±0.05	0.9±0.06	0.95±0.04
tfidf	naiveBayes	0.76±0.09	0.76±0.08	0.76±0.1

Table 1: Cross-validation accuracy with removal of stop words

computationally expensive, and the scarcity of advertorials limited the use of an additional hold out test set. This could be addressed in future work.

For SVM, SGD and linearSVC we increased the maximum amount of iterations to 5000 and for decision tree and random forest we set the max depth to “none”. In terms of evaluation we ran 10-fold cross validation to test various algorithms and parameters, as well as a cross domain test set up where one medium is used as the test set, and models are trained on the other media. The metrics that we evaluate our results with are accuracy, f1 score and AUC.

4.2 Results

In a first set of experiments we benchmark the performance of all algorithms across regular and TF-IDF weighted BoW representations. Table 1 shows the results, with stop words removed; the results with stop words included were very similar. TF-IDF typically outperformed regular BoW so the remainder of the experiments was carried out with TF-IDF, with stop word filtering.

The results of the cross domain testing experiment can be found in Table 2. The best results were obtained with SVM, linearSVC, random forest and SGD, closely followed by decision trees and naive bayes, and k-NN scored poorly, probably due to high dimensionality. Top scoring results were close, but SVM scored best, so we decided to continue the experiments with this method. In terms of media, NRC scored best, followed by Nu.nl and Telegraaf, and Ondernemer scoring substantially worse, which may be due to the fact that in the business to business domains editorial and commercial content is more similar.

We also ran a structured experiment where we gradually increased the number of features that made clear that at 5000 features performance more or less stabilizes (results omitted for brevity), and we ran a series of tests to study the impact of tweaking the various SVM parameters (Table 3).

	SVM	Linear SVC	Decision Tree	Random Forest	k-NN	SGD	Naive Bayes	
Nu.nl	0.84	0.84	0.72	0.83	0.52	0.84	0.72	0.76 ± 0.12
NRC	0.93	0.95	0.75	0.82	0.52	0.95	0.81	0.82 ± 0.15
Ondernemer	0.76	0.68	0.54	0.55	0.51	0.65	0.56	0.61 ± 0.09
Telegraaf	0.85	0.84	0.66	0.84	0.46	0.83	0.73	0.74 ± 0.14
	0.85 ± 0.06	0.83 ± 0.10	0.67 ± 0.08	0.76 ± 0.12	0.5 ± 0.02	0.82 ± 0.11	0.71 ± 0.09	

Table 2: Cross-domain testing results (test set in rows, trained on other media, metric is accuracy)

To further validate the cross domain results, we also trained and tested models with data from just one medium each, and created t-SNE graphs (Figure 2, along with the corresponding accuracies). t-SNE graphs [9] are a way to represent multi-dimensional data (in our case a 5000 dimensions) in a two-dimensional scatter plot. For our experiment, we ran the t-SNE graph with a perplexity of 30, a maximum number of iterations of 1000 and a random state of 2. In other words, apart from the random state only the default parameter values.

Using this method we can visualize how well the classes can be separated based on the available data, making it possible to visualize the separation of church and state in our experiment. The ranking of various media are consistent with the cross domain results, with NRC displaying the clearest separation and Ondernemer the worst.

After completing the experimental process explained earlier we found that the model explained in Table 4 gave us the best results (all other parameters are defaults). So we used this model to derive a lexicon by training a model on all data and using this model’s feature terms and weights. This is useful, even though it serves the same purpose as our model, because it can be published without publishing the data as well, which we are not able to do because of copyright issues, and it can be consumed more easily by a broad non technical audience such as journalists and social scientists.⁵

Using a linear kernel means that the separating hyperplane is defined in the original input space, hence we can interpret the weights of the model as term weights in a lexicon. Users can make very simple use of the lexicon, just by counting the occurrence of negative and positive words (with zero as threshold) or approximate the original model closer, for example by calculating a score by multiplying frequency of the terms with the term weights and summing the results. Figure 3 shows the distribution of the scores for the full corpus, calculated with the latter approach. One can clearly see two more or less normal distributions representing the advertorials and regular articles.

Inspection of these feature coefficients also provides further insight into differences in language use between classes. In Figure 4 we have listed the features

⁵ The lexicon is published at https://github.com/TimoKats/research_distinguishing_commercial_and_editorial_content

kernel	decision function	accuracy	f1_score	roc_auc
linear	ovo	0.9029±0.0559	0.9003±0.0581	0.9495±0.0341
linear	ovr	0.9029±0.0559	0.9003±0.0581	0.9495±0.0341
poly	ovo	0.8094±0.0774	0.7755±0.1016	0.9167±0.0412
poly	ovr	0.8094±0.0774	0.7755±0.1016	0.9167±0.0412
rbf	ovo	0.8999±0.0579	0.8989±0.0583	0.9438±0.0393
rbf	ovr	0.8999±0.0579	0.8989±0.0583	0.9438±0.0393
sigmoid	ovo	0.9009±0.0563	0.8986±0.0585	0.9498±0.0339
sigmoid	ovr	0.9009±0.0563	0.8986±0.0585	0.9498±0.0339

Table 3: The effect of tweaking the parameters with svm

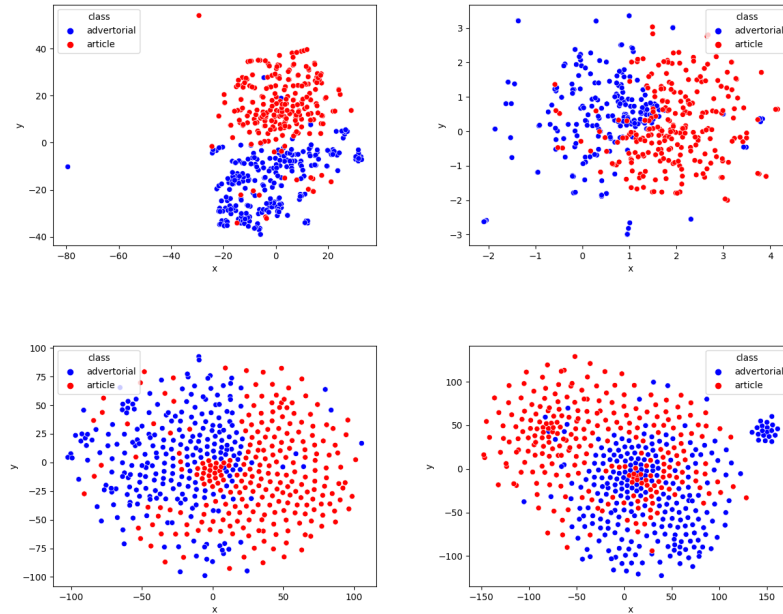


Fig. 2: t-SNE plots for NRC (95%), Nu.nl (92%), Telegraaf (91%) and Onderneer (85%)

learning model	features	text representation	kernel	max iter	accuracy
svm	5000	tf-idf	linear	5000	0.9029±0.0559

Table 4: Settings from the final model.




Fig. 3: The distribution of the predicted scores from the entries using the lexicon.




Fig. 4: Coefficients with highest absolute values for regular articles (left) and advertorials (right)

with the highest absolute values for regular articles and advertorials. As can be seen, the difference isn't just a matter of topics (f.i. 'cabinet', 'minister' for news, 'investing', 'enterprise', 'technology', 'innovation' for advertorials), but also a matter of how these topics are being talked about. In regular articles, indications of time (days, months etc) and attribution ('writes', 'says', 'appeared') score high, whereas high scoring features for advertorials include adjectives such as 'free', 'healthy' and 'sustainable', perhaps highlighting the benefits of products and services. Coefficients for the full 5000 features can be downloaded with the lexicon, and we also investigated by training models on media separately to understand differences between publications.

5 Exploring the Corpus with Co-occurrence Networks

Results such as feature importance already provided us with some insights into how language differs between advertorials and regular articles, but to delve




Fig. 5: Overview of the co-occurrence network. Complete version can be found at: timokats.github.io/network/

deeper in a more exploratory fashion we have also created co-occurrence networks (see Figure 5).

The nodes in this network are the terms from the lexicon, blue nodes are the editorial terms and the red nodes are the commercial terms (negative and positive coefficients). The size of the node is related to its degree. The edges represent the connection between two terms in the data set. We calculated this based on how often the two terms appear in the same sentence as a percentage. So for example, in our data set every time the term ‘artificial’ appears in a sentence, 75% of the time that sentence also has the word ‘intelligence’. Thus, there’s a directed edge from ‘artificial’ to ‘intelligence’. For visualization we show all nodes with edges exceeding a minimum threshold, or likewise this could be seen as an undirected graph where the weight is the lowest value for the two terms.

By exploring the co-occurrence network certain things about both our results and data become apparent. First, the fact that some of our data centers around subjects that were very prevalent in 2020 (like the US elections and the covid-19 pandemic), resulting in a time frame bias towards 2020, because in fu-

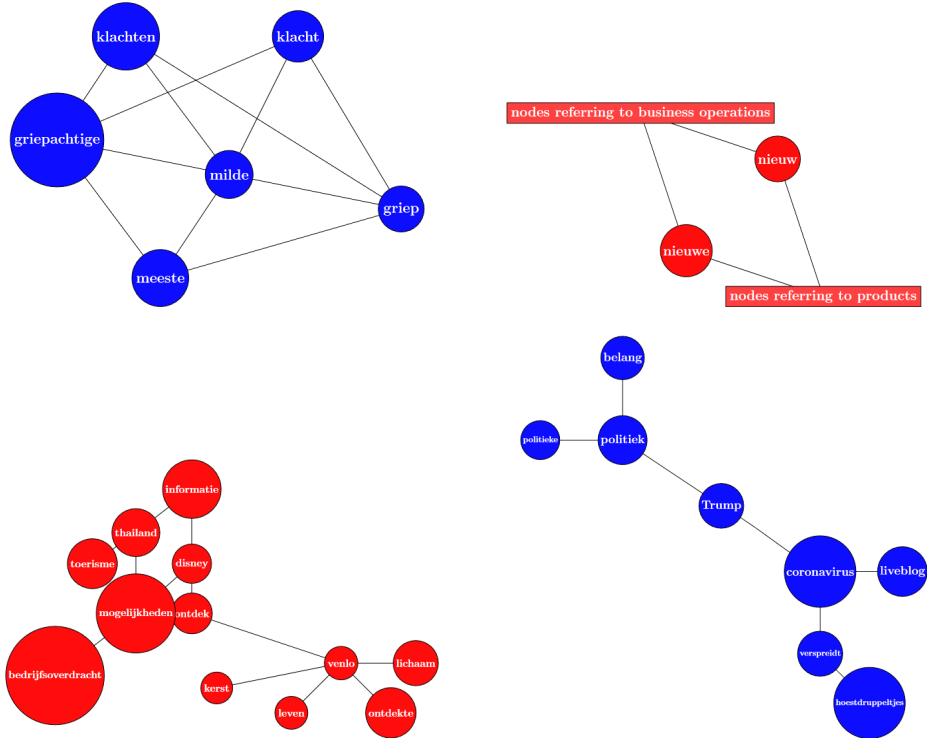


Fig. 6: Commercial and editorial sub graphs.

ture implementations of our model and/or lexicon these subjects may be less prevalent.

Second, it has also given us more insight into the structure of commercial language and how it’s different from editorial language. For example, commercial language in our network has two large clusters (one related to goods and one related services). These clusters are linked by the terms ‘nieuw’ (new) and ‘nieuwe’ (new). For our editorial clusters we for example found a cluster related to covid symptoms, which showcases the time frame bias mentioned earlier. Through using a co-occurrence graph we can find patterns/clusters like these and gain more insight into our data and results. An overview of some important findings in our graph can be found in Figure 6.

6 Conclusion

This research aims to differentiate commercial and editorial content, and more specifically, advertorials from regular articles, and our main research questions are the following. To what extent can we differentiate advertorials and articles by using machine learning? And can we use machine learning and a data driven

approach to better understand the difference between commercial and editorial language?

We answered the first question by developing a range of models for various media, and deriving a lexicon from it. The best models perform with over 90 per cent accuracy, and as mentioned this is an optimistic estimate and performance clearly varies by medium and set up. Further insight is provided by highlighting the differences of performance across media, with business-to-business medium *Ondernemer* scoring lowest, which could make sense given similarities in jargon. Feature importance analysis and co-occurrence graphs provided further insight into differences in language, both from a topic perspective, as well as how these topics were being spoken about.

Our research has some known limitations. In particular the size of the data set (of just 2000 entries) could be increased in future work, including a wider set of media and longer time frames. A key challenge here to overcome is that that particularly advertorials are not always available for extended periods of time. It may also be interesting to expand the scope to other major languages and other forms of native advertising. We also plan to engage with the general public, journalists as well as marketers, using the results of this research to raise awareness and trigger debate and discussion.

Despite some of its limitations we think our research can serve as an example to put the problem on the agenda, provide insight into it, and illustrate the potential of using machine learning for differentiating commercial and editorial content. Moreover, it also showcases how machine learning and AI can be a solution, not just a problem, in society and the modern digital media landscape.

References

1. Cameron, G.T., Ju-Pak, K.H., Kim, B.H.: Advertorials in magazines: Current use and compliance with industry guidelines. *Journalism & Mass Communication Quarterly* **73**(3), 722–733 (1996)
2. Conill, R.F.: Camouflaging church as state: An exploratory study of journalism’s native advertising. *Journalism Studies* **17**(7), 904–914 (2016)
3. Cornia, A., Sehl, A., Nielsen, R.K.: ‘We no longer live in a time of separation’: A comparative analysis of how editorial and commercial integration became a norm. *Journalism* **21**(2), 172–190 (2020)
4. Graham, J., Haidt, J., Nosek, B.A.: Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology* **96**(5), 1029 (2009)
5. Iversen, M.H., Knudsen, E.: When politicians go native: The consequences of political native advertising for citizens’ trust in news. *Journalism* **20**(7), 961–978 (2019)
6. Kats, T.: Differentiating commercial and editorial content (2021), BSc thesis, Leiden University
7. Kim, B.H., Pasadeos, Y., Barban, A.: On the deceptive effectiveness of labeled and unlabeled advertorial formats. *Mass Communication and Society* **4**(3), 265–281 (2001)

8. Krouwer, S., Poels, K., Paulussen, S.: To disguise or to disclose? the influence of disclosure recognition and brand presence on readers' responses toward native advertisements in online news media. *Journal of Interactive Advertising* **17**, 00–00 (10 2017)
9. van der Maaten, L., Hinton, G.: Visualizing data using t-SNE. *Journal of Machine Learning Research* **9**, 2579–2605 (11 2008)
10. Pew Research Center: State of the media: Newspapers. <https://www.pewresearch.org/journalism/fact-sheet/newspapers/>, Online: accessed: 2021-09-07
11. van Reijmersdal, E., Neijens, P., Smit, E.: Readers' reactions to mixtures of advertising and editorial content in magazines. *Journal of Current Issues & Research in Advertising* **27**(2), 39–53 (2005)
12. Schelling, J., van Eekelen, N., van Veelen, I., van Hees, M., van der Putten, P.: Bursting the bubble (extended abstract). In: MISDOOM 2020. p. 72 (10 2020)
13. Taboada, M., Brooke, J., Tofiloski, M., Voll, K., Stede, M.: Lexicon-based methods for sentiment analysis. *Computational Linguistics* **37**, 267–307 (06 2011)
14. Teernstra, L., van der Putten, P., Noordegraaf-Eelens, L., Verbeek, F.J.: The morality machine: Tracking moral values in tweets. In: Boström, H., Knobbe, A.J., Soares, C., Papapetrou, P. (eds.) *Advances in Intelligent Data Analysis XV - 15th International Symposium, IDA 2016, Stockholm, Sweden, October 13-15, 2016, Proceedings. Lecture Notes in Computer Science*, vol. 9897, pp. 26–37 (2016)
15. Wilkinson, J.B., Hausknecht, D.R., Prough, G.E.: Reader categorization of a controversial communication: Advertisement versus editorial. *Journal of Public Policy & Marketing* **14**(2), 245–254 (1995)
16. Wojdowski, B., Evans, N.: Going native: Effects of disclosure position and language on the recognition and evaluation of online native advertising. *Journal of Advertising* pp. 1–12 (2016)
17. Zhou, S.: 'Advertorials': A genre-based analysis of an emerging hybridized genre. *Discourse & Communication* **6**(3), 323–346 (2012)