

# MN-DS: A Multilabeled News Dataset for News Article Hierarchical Classification

Alina Petukhova \*  and Nuno Fachada 

COPELABS, Lusófona University, Campo Grande 376, 1749-024 Lisbon, Portugal

\* Correspondence: alina.petukhova@ulusofona.pt

**Abstract:** This article presents a dataset of 10,917 news articles with hierarchical news categories collected between 1 January 2019 and 31 December 2019. We manually labeled the articles based on a hierarchical taxonomy with 17 first-level and 109 s-level categories. This dataset can be used to train machine learning models for automatically classifying news articles by topic. This dataset can be helpful for researchers working on news structuring, classification, and predicting future events based on released news.

**Keywords:** news dataset; text classification; NLP; media topic taxonomy

## 1. Background and Summary

A news dataset is a collection of news articles classified into different categories. In the past decade, there has been a sharp increase in news datasets available for analysis [1]. These datasets can be used to understand various topics, from politics to the economy.

A few different types of news datasets are commonly used for analysis. The first is raw data, which includes all the data that a news organization collects. This data can be used to understand how a news organization operates, what stories are covered, and how they are covered. The second type of news dataset is processed data. These data have been through some processing, such as aggregation or cleaned up. Processed data are often easier to work with than raw data and can be used to answer specific questions such as providing additional information for the decision-making process. The third type of news dataset is derived data. These data are created by combining multiple datasets, often from different sources [2]. News datasets can be used for various purposes in a machine learning context, for example:

- Predicting future events based on past news articles.
- Understanding the news cycle.
- Determining the sentiment of news articles.
- Extracting information from news articles (e.g., named entities, location, dates).
- Classifying news articles into predefined categories.

To adequately answer research questions, news datasets should contain sufficient data points and span a significant enough period. There are many labeled news datasets available, each with specific limitations. For example, they may only cover a specific period or geographical area or be confined to a particular topic. Additionally, the categories may not be completely accurate, and the datasets may be biased in some way [3,4].

Some of the more popular news datasets include the 20 Newsgroups dataset [5], AG's news topic classification dataset [6], L33-Yahoo News dataset [7,8], News Category dataset [9], and Media Cloud dataset [10]. Each of these datasets has been used extensively by researchers in the fields of natural language processing and machine learning, and each has its advantages and disadvantages. The 20 Newsgroups dataset was created in 1997 and contains 20 different categories of news, each with a training and test set. The data



**Citation:** Petukhova, A.; Fachada, N. MN-DS: A Multilabeled News Dataset for News Article Hierarchical Classification. *Data* **2023**, *8*, 74. <https://doi.org/10.3390/data8050074>

Academic Editor: Kassim S. Mwitondi

Received: 19 March 2023

Revised: 17 April 2023

Accepted: 19 April 2023

Published: 23 April 2023



**Copyright:** © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

is already pre-processed and tokenized, which makes it very easy to use. However, the dataset is outdated and relatively small, with only about 1000 documents in each category.

The AG's news topic classification dataset is a collection of news articles from the academic news search engine "ComeToMyHead" during more than one year of activity. Articles were classified into 13 categories: business, entertainment, Europe, health, Italia, music feeds, sci/tech, software & dev., sports, toons, top news, U.S., and world. The dataset contains more than 1 million news articles. However, there are several limitations to this dataset. First, it is currently outdated since data were collected in 2005. Second, the taxonomy covers specific countries such as the US and Italy but has general references such as Europe or world, creating overlaps in the classification (e.g., Italy and Europe) as well as potential imbalances (e.g., events in China are likely to be underrepresented and/or under-reported compared to those in the US). Finally, the dataset does not include methods for type or category description.

The L33-Yahoo News dataset is a collection of news articles from the Yahoo News website provided as part of the Yahoo! Webscope program. The articles are labeled into 414 categories such as music, movies, crime justice, and others. The dataset includes the random article id followed by possible associated categories. The L33-Yahoo News dataset is available under Yahoo's data protection standards. It can be used for non-commercial purposes if researchers credit the source and license new creations under identical terms. The limitations of the L33 dataset are the license terms, restricting companies from using this dataset for commercial purposes, and the amount of data per class, with the category "supreme court decisions" having only five articles, for example. In addition, there is some overlap in the categories, which makes it challenging to train a model that can accurately predict multiple categories.

The News Category Dataset is a collection of around 210k news articles from the Huffington Post, labeled with their respective categories, which include business, entertainment, politics, science and technology, and sports. However, the dataset has several limitations. First, the dataset is not comprehensive since it only includes articles from one source. Second, news categories are not standardized, including broad categories such as "Media" and "Politics" and very narrow ones like "Weddings" and "Latino voices".

The Media Cloud Data Set is a collection of over 1.7 billion articles from more than 60 thousand media sources around the world. The dataset includes articles from both mainstream and alternative news sources, including newspapers, magazines, blogs, and online news outlets. Data can be queried by keyword, tag, category, sentiment, and location. This dataset is useful for researchers who are interested in studying media coverage of specific topics or trends over time. Media Cloud is a large multilingual dataset that has good media coverage but limited use in topic classification models since it does not include a mapping of articles to a specific news taxonomy.

The main motivation for this work is to provide a dataset for building specific topic models. It consists of a categorized subset taken from an existing news dataset. We show that such a dataset, with up-to-date articles mapped into a standardized news taxonomy, can contribute to the accuracy improvement of news classification models.

## 2. Methods

In this paper, we present a new dataset based on the NELA-GT-2019 data source [11], classified with IPTC's NewsCodes Media Topic taxonomy [12] (The International Press Telecommunications Council, or IPTC, is an organization that creates and maintains standards for exchanging news and other information between news organizations). The original NELA-GT-2019 dataset contains 1.12 M news articles from 260 sources collected between 1 January 2019 and 31 December 2019, providing essential content diversity and topic coverage. Sources include a wide range of mainstream and alternative news outlets.

In turn, the IPTC taxonomies are a set of controlled vocabularies used to describe news stories' content. The NewsCodes Media Topic taxonomy has been one of IPTC's main subject taxonomies for text classification since 2010. We used the 2020 version of

NewsCodes Media Topic taxonomy [13]. News organizations use it to categorize and index their content, while search engines use it to improve the discoverability of news stories [14].

Algorithm of the article selection process:

1. Obtain a random article from the NELA dataset;
2. Classify it for the second-level category of the NewsCodes Media Topic taxonomy by checking the keywords and thorough reading of the article; the news article is assigned to exactly one category;
3. If there are already 100 articles in that category discard it, otherwise assign a second-level category to the article;
4. Return to step 1 and repeat until each second-level category has 100 articles assigned.

The described algorithm allows for overcoming the limitation of the NELA-GT datasets where a large proportion of the dataset is fringe, conspiracy-based news due to the discharging of the news if a category already has 100 articles in it.

We observed that the first-level category of the NewsCodes Media Topic taxonomy is not accurate enough to catalogue an article. For example, the “sport” category may include different aspects, such as information about specific sports, sports event announcements, and the sports industry in general, which have more specific meanings than the first-level category label is able to convey. Therefore, we used a second-level category of NewsCodes Media Topic taxonomy to have a more specific article category. In comparison to the previously published datasets, we included in our dataset unique categories such as “arts and entertainment”, “mass media”, “armed conflict”, “weather statistic”, and “weather warning”. Therefore, we created the proposed Multilabeled News Dataset (MN-DS) by hand-picking and labeling approximately 100 news articles for each second level category (<https://www.iptc.org/std/NewsCodes/treeview/mediatopic/mediatopic-en-GB.html> (accessed on 13 March 2022)) of the NewsCodes Media Topic taxonomy.

### 3. Data Records

After manually selecting news articles relevant to each category, we obtained 10,917 articles in 17 first-level and 109 second-level categories from 215 media sources. During the selection process, one article was processed by one coder. An overview of the released MN-DS dataset by category is provided in Table 1. All data are available in CSV format at <https://doi.org/10.5281/zenodo.7394850> under a Creative Commons license.

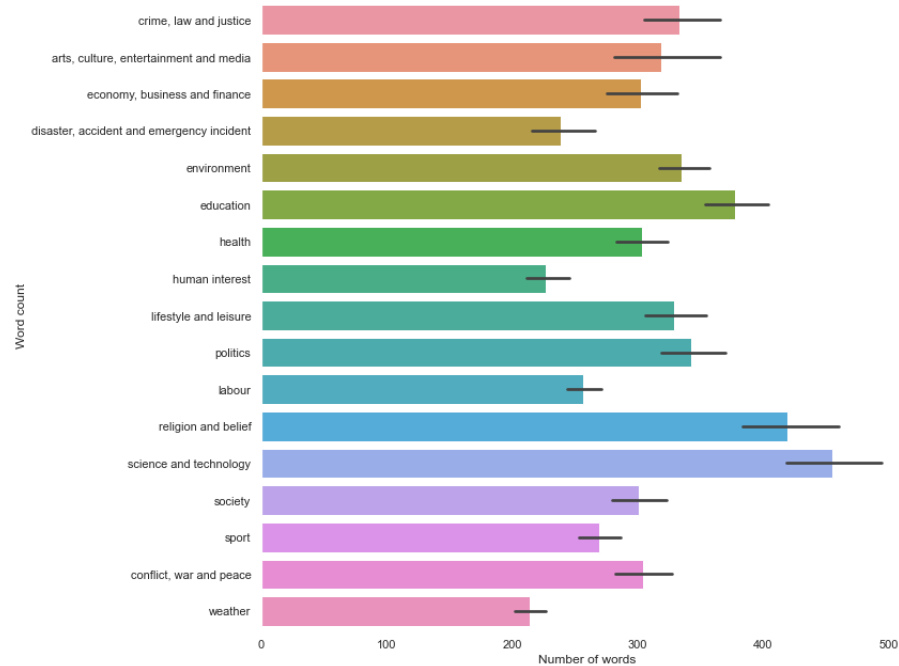
**Table 1.** The number of articles under each Level 1 category.

Categories	Count
Arts, culture, entertainment, and media	300
Conflict, war, and peace	800
Crime, law, and justice	500
Disaster, accidents, and emergency incidents	500
Economy, business and finance	400
Education	607
Environment	600
Health	700
Human interest	600
Labor	703
Lifestyle and leisure	300
Politics	900
Religion and belief	800
Science and technology	800
Society	1100
Sport	907
Weather	400

The MN-DS contains articles published in 2019, the distribution of selected articles over the year is balanced with slightly more articles for the month of January 2019. The

To describe the dataset, we created a word cloud representation of each category, as shown in Figure 1. The central concept of a word cloud is to visualize for each category the most popular words with a size corresponding to the degree of popularity. This representation allows us to quickly assess the quality of the text annotation since it displays the most common words of the category. In the bar chart shown in Figure 2, we can observe that the “science and technology” first-level category contains the highest count of topic-specific words, while in more general categories, such as “weather” or “human interest”, there is less variety in the texts, probably because they represent shorter and more similar articles.

- 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58, 59, 60, 61, 62, 63, 64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 74, 75, 76, 77, 78, 79, 80, 81, 82, 83, 84, 85, 86, 87, 88, 89, 90, 91, 92, 93, 94, 95, 96, 97, 98, 99



**Figure 2.** Mean number of non-repeated words in article body for first-level categories. The error bars represent the 95% confidence interval.

#### 4. Usage Example

We used the dataset to train the most common text classification models to extend the technical validation of the proposed dataset and establish the benchmark for multiclass classification. The following embeddings were selected:

- Tf-idf embedding, where Tf-idf stands for term frequency-inverse document frequency [15]. Tf-idf transforms text into a numerical representation called a tf-idf matrix. The term frequency is the number of times a word appears in a document. The inverse document frequency measures how common a word is across all documents. Tf-idf is used to weigh words so that important words are given more weight. The dataset's news texts and categories were combined and vectorized with TfidfVectorizer [16].
- GloVe (Global Vectors for Word Representation) embeddings with an algorithm based on a co-occurrence matrix, which counts how often words appear together in a text corpus. The resulting vectors are then transformed into a lower-dimensional space using singular value decomposition [17].
- DistilBertTokenizer [18], which is a distilled version of BERT, a popular pre-trained model for natural language processing. DistilBERT is smaller and faster than BERT, making it more suitable for fast training with limited resources. The trade-off is that DistilBERT's performance is 3% lower than BERT's. DistilBERT embeddings are trained on the same data as BERT, so they are equally good at capturing the meaning of words in context.

During dataset validation, we combined the selected embeddings with different classifiers. We tested multinomial naive Bayes (NB) classifier [19], logistic regression [20], support vector classifier (SVC) [21], and DistilBERT model [22]. Since MN-DS is a multiclass dataset, we used the OneVsRestClassifier strategy for classification models [16]. OneVsRestClassifier is a classifier that trains multiple binary classifiers, one for each class. The individual binary classifiers are then combined to create a single multiclass classifier. This approach is often used when there are many categories, as it can be more efficient than training a single multiclass classifier from scratch. The tested classifiers work as follows:

- The multinomial NB is a text classification algorithm that uses Bayesian inference to classify text. It is a simple and effective technique that can be used for various tasks, such as spam filtering and document classification. The algorithm is based on the assumption that the features in a document are independent of each other, which allows it to make predictions about the category of a document based on its individual features.
- The logistic regression classifier works by using a sigmoid function to map data points from an input space to an output space, where the categories are assigned based on a linear combination of the features. The weights of the features are learned through training, and the predictions are made by taking the dot product of the feature vector and the weight vector.
- The SVC classifier is a powerful machine learning model based on the support vector machines algorithm. The model is based on finding the optimal decision boundary between categories to maximize the margin of separation between them. The SVC model can be used for linear and non-linear classification tasks and is particularly well-suited for problems with high dimensional data. The classifier is also robust to overfitting and can generalize well to new data.
- DistilBERTModel, a light version of the BERT classifier [18], developed and open-sourced by the team at Hugging Face. DistilBERTModel can be fine-tuned with just one additional output layer to create state-of-the-art models for a wide range of NLP tasks with minimal training data.

Classification results for level 1 and level 2 categories are presented in Tables 2 and 3, respectively. It is possible to observe that DistilBERTModel achieves better classification results for both category levels. To improve these results in future studies, we suggest applying hierarchical classification methods as described by Silla and Freitas [23], for example.

**Table 2.** Multilabel classification results for level 1 categories.

Embeddings Model	TFIDF			Glove			DistilBertTokenizer		
	Precision	Recall	f1 Score	Precision	Recall	f1 Score	Precision	Recall	f1 Score
Multinomial NB	0.802	0.631	0.649	0.629	0.499	0.529	n/a	n/a	n/a
Logistic Regression	0.800	0.763	0.774	0.747	0.739	0.739	n/a	n/a	n/a
SVC Classifier	0.808	0.796	0.799	0.768	0.762	0.760	n/a	n/a	n/a
DistilBERTModel	n/a	n/a	n/a	n/a	n/a	n/a	0.849	0.842	0.844

**Table 3.** Multilabel classification results for level 2 categories.

Embeddings Model	TFIDF			Glove			DistilBertTokenizer		
	Precision	Recall	f1 Score	Precision	Recall	f1 Score	Precision	Recall	f1 Score
Multinomial NB	0.628	0.602	0.583	0.496	0.484	0.469	n/a	n/a	n/a
Logistic Regression	0.646	0.649	0.635	0.589	0.589	0.577	n/a	n/a	n/a
SVC Classifier	0.645	0.646	0.628	0.581	0.595	0.571	n/a	n/a	n/a
DistilBERTModel	n/a	n/a	n/a	n/a	n/a	n/a	0.735	0.715	0.715

**Author Contributions:** A.P. created the search strategy, retrieved and screened the publications, extracted the selected data for annotation, and labeled the dataset. A.P. and N.F. assessed the quality of the included articles and checked the data. A.P. performed the statistical analyses, created the graphics and wrote the original draft. N.F. conceived the project, provided critical comments, and revised the paper. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research was funded by Fundação para a Ciência e Tecnologia under Project UIDB/04111/2020 (COPELABS).

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** The data described in this paper is available in CSV format at <https://doi.org/10.5281/zenodo.7394850>. Code for the technical validation of the dataset is available at <https://github.com/alinaPETUKHOVA/mn-ds-news-classification> (accessed on 15 April 2023).

**Conflicts of Interest:** The authors declare no conflict of interest.

## References

1. Paullada, A.; Raji, I.D.; Bender, E.M.; Denton, E.; Hanna, A. Data and its (dis)contents: A survey of dataset development and use in machine learning research. *Patterns* **2021**, *2*, 100336. [[CrossRef](#)] [[PubMed](#)]
2. Jayakody, N.; Mohammad, A.; Halgamuge, M. Fake News Detection using a Decentralized Deep Learning Model and Federated Learning. In Proceedings of the IECON 2022—48th Annual Conference of the IEEE Industrial Electronics Society, Brussels, Belgium, 17–20 October 2022.
3. Stefansson, J.K. Quantitative Measure of Evaluative Labeling in News Reports: Psychology of Communication Bias Studied by Content Analysis and Semantic Differential. Master's Thesis, UiT, Norway's Arctic University, Tromsø, Norway, 2014.
4. Gezici, G. Quantifying Political Bias in News Articles. *arXiv* **2022**, arXiv:2210.03404.
5. Mitchell, T. 20 Newsgroups Data Set. 1999. Available online: <http://qwone.com/~jason/20Newsgroups/> (accessed on 10 April 2023).
6. AG's Corpus of News Articles. 2005. Available online: [http://groups.di.unipi.it/~gulli/AG\\_corpus\\_of\\_news\\_articles.html](http://groups.di.unipi.it/~gulli/AG_corpus_of_news_articles.html) (accessed on 10 April 2023).
7. Soni, A.; Mehdad, Y. RIPML: A Restricted Isometry Property-Based Approach to Multilabel Learning. In Proceedings of the Thirtieth International Florida Artificial Intelligence Research Society Conference, FLAIRS 2017, Marco Island, FL, USA, 22–24 May 2017; Rus, V., Markov, Z., Eds.; AAAI Press: Palo Alto, CA, USA, 2017; pp. 532–537.
8. Chen, S.; Soni, A.; Pappu, A.; Mehdad, Y. DocTag2Vec: An Embedding Based Multi-label Learning Approach for Document Tagging. In Proceedings of the Rep4NLP@ACL, Vancouver, BC, Canada, 3 August 2017.
9. Misra, R. News Category Dataset. *arXiv* **2022**, arXiv:2209.11429.
10. Roberts, H.; Bhargava, R.; Valiukas, L.; Jen, D.; Malik, M.; Bishop, C.; Ndulue, E.; Dave, A.; Clark, J.; Etling, B.; et al. Media Cloud: Massive Open Source Collection of Global News on the Open Web. *arXiv* **2021**, arXiv:2105.18127.
11. Gruppi, M.; Horne, B.D.; Adalı, S. NELA-GT-2019: A Large Multi-Labelled News Dataset for The Study of Misinformation in News Articles. *arXiv* **2020**, arXiv:2003.08444.
12. IPTC NewsCodes Scheme (Controlled Vocabulary). 2010. Available online: <https://cv.iptc.org/newscodes/mediatopic/> (accessed on 10 April 2023).
13. IPTC Media Topics—Vocabulary Published on 25 February 2020. 2020. Available online: [https://www.iptc.org/std/NewsCodes/previous-versions/IPTC-MediaTopic-NewsCodes\\_2020-02-25.xlsx](https://www.iptc.org/std/NewsCodes/previous-versions/IPTC-MediaTopic-NewsCodes_2020-02-25.xlsx) (accessed on 10 April 2023).
14. NewsCodes—Controlled Vocabularies for the Media. Available online: <https://iptc.org/standards/newscodes/#:~:text=Who%20uses%20IPTC%20NewsCodes%3F,becoming%20more%20and%20more%20popular> (accessed on 21 November 2022).
15. Sammut, C.; Webb, G.I. (Eds.) TF-IDF. In *Encyclopedia of Machine Learning*; Springer: Boston, MA, USA, 2010; pp. 986–987. [[CrossRef](#)]
16. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-Learn: Machine Learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830.
17. Pennington, J.; Socher, R.; Manning, C. Glove: Global Vectors for Word Representation. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), Doha, Qatar, 25–29 October 2014; pp. 1532–1543. [[CrossRef](#)]
18. Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al. Transformers: State-of-the-Art Natural Language Processing. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, Online, 16–20 November 2020; pp. 38–45.
19. Manning, C.D.; Raghavan, P.; Schütze, H. *Introduction to Information Retrieval*; Cambridge University Press: Cambridge, MA, USA, 2008; pp. 234–265.
20. Cox, D.R. The regression analysis of binary sequences. *J. R. Stat. Soc. Ser. B (Methodol.)* **1958**, *20*, 215–242. [[CrossRef](#)]
21. Cortes, C.; Vapnik, V. Support-vector networks. *Mach. Learn.* **1995**, *20*, 273–297. [[CrossRef](#)]
22. Sanh, V.; Debut, L.; Chaumond, J.; Wolf, T. DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter. *arXiv* **2019**, arXiv:1910.01108.
23. Silla, C.; Freitas, A. A survey of hierarchical classification across different application domains. *Data Min. Knowl. Discov.* **2011**, *22*, 31–72. [[CrossRef](#)]

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.