

Exploring Sentiment Analysis in Health with CNN, GRU and Zero-Shot Learning on Imbalanced Textual Data.

Souaad HAMZA-CHERIF¹, Nesma SETTOUTI^{1,2}

¹Biomedical Engineering Laboratory, University of Tlemcen, Algeria

²L@bISEN, Yncréa ouest, Brest, France

Abstract

Sentiment analysis has recently gained popularity due to the emergence of social media and the semantic web, which provide platforms for individuals to express their thoughts and emotions. In the field of health, analyzing this type of data can have numerous benefits, particularly in psychology, where it can aid in automatically identifying a patient's mental state. In this study, we present a method for categorizing emotions from text data in the health field. We evaluated various classifiers such as: Convolutional Neural Networks (CNN), Gated Recurrent Units (GRU), and Zero-Shot learning (ZSL) on the unbalanced EmoHD dataset. Our findings showed that the CNN model gives the best performance[1].

Keywords

Sentiment analysis, CNN, GRU, Zero-Shot learning, classification.

1. Introduction

Sentiments are complex phenomena that can take many forms, such as emotions, reactions to events, or mental states. They can be positive or negative, joyful or boring, sad, etc.

Sentiment analysis has become a rapidly growing field, particularly with the rise of social media and the semantic web, which make it easier for people to express their thoughts and emotions through text, emoticons, and other means.

Machine understanding and analysis of sentiments is particularly important in the field of health, where a patient's emotional state can greatly impact their healing process. For example, a patient suffering from depression may not be in a favorable mental state to begin their recovery. Similarly, in the field of psychology, automatic sentiment analysis can be very beneficial in identifying a patient's psychological state as well as the different human personality traits that can explain human reasoning and behavior.

In this context, artificial learning techniques and natural language processing (NLP) offer promising tools which have proven themselves to provide automatic solutions in different fields such as health because they play an important role in understanding the meaning of content for classifying, analyzing, and predicting human sentiments by machine. In this article, we present

RIF 2023 : The 12th Seminary of Computer Science Research at Feminine, March 9, 2023, Constantine, Algeria

✉ souad.hamzacherif@univ-tlemcen.dz (S. HAMZA-CHERIF); nesma.settouti@univ-tlemcen.dz (N. SETTOUTI)

🆔 0000-0002-4733-197X (S. HAMZA-CHERIF); 0000-0002-7423-0090 (N. SETTOUTI)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

an approach for sentiment analysis from text data in the field of health. We compare various machine learning methods, including convolutional neural networks (CNN), gated recurrent units (GRU), and zero-shot learning (ZSL).

The rest of the article is structured as follows. In section 2, we will discuss related work in the field, then we will present our proposed approach and research strategy in section 3. In section 4, we will analyze the results we obtained and conclude with perspectives and future works.

2. Related works

Several recent studies have focused on using artificial learning for sentiment analysis, one notable example is in [2], where the authors used a lexicon-based classification approach to predict the results of the 2016 U.S. Presidential Election between Hillary Clinton and Donald Trump using Twitter data.

In [3], the authors analyzed sentiments towards AstraZeneca, Pfizer and Moderna COVID-19 vaccines using Twitter data and the AFINN lexicon. The results showed positive sentiment towards Pfizer and Moderna vaccines. In [4] the authors used the Word2vec model to compute vector representations of words and then applied the Convolutional Neural Network CNN model to analyze sentiment from a corpus of movie review excerpts that includes five labels (negative, rather negative, neutral rather positive and positive) and achieved a test accuracy of 45.4%.

Other studies have also applied sentiment analysis to health-related textual data. In [5], the authors used term frequency-inverse document frequency (TF-IDF) vectorization and compared the performance of different classifiers for emotion classification on the EmoHD dataset. In [6], the authors used the Intelligent Water Drop algorithm to select features of interest from the EmoHD dataset and classify emotions in the health domain.

In [7], the authors used deep learning models for Arabic aspect-based sentiment analysis. To take advantage of both word and character representations, and extract the main opinion aspect, they combined a bidirectional GRU, a Convolutional Neural Network (CNN), and a Conditional Random Field (CRF), and then used an interactive attention network based on a bidirectional GRU (IAN-BGRU) to identify the sentiment polarity towards the extracted aspects.

In [8], the authors trained a unified model that performs zero-shot aspect-based sentiment analysis (ABSA) without using any annotated data for a new domain using Zero-Shot Learning. They evaluated it for end-to-end aspect-based sentiment analysis (E2E ABSA), which shows that ABSA can be conducted without any human-annotated ABSA data.

In [9] the authors also use the modified GRU model for the classification and understanding of sentiments from tweets from an online site and compared it to the long-short-term memory model (LSTM) and the bidirectional long-short-term memoryterm (BiLSTM) and demonstrated that the modified GRU outperformed them. In [10], the authors propose a two-stage emotion detection methodology. In the first stage, they use a Zero-Shot Learning model based on a sentence transformer, returning the probabilities for subsets of 34 emotions. Then, they used the output of the zero-shot model as an input for the second stage, which trains a machine learning classifier on the sentiment labels in a supervised manner using ensemble learning.

In this work, we compare different approaches for classifying sentiments from health-related textual data (EmoHD) using CNN, GRU, and ZSL models. Our primary goal is to explore other deep learning models on the EmoHD dataset and evaluate their performance especially since these models have demonstrated their effectiveness in the field. Then, we aim to evaluate the annotation of EmoHD dataset using the Zero-Shot Learning model.

3. Proposed approach

In this article, we present a method for classifying sentiments from text data in the EmoHD [5] health dataset. The EmoHD dataset is composed of 4,202 text samples from eight disease classes and six emotion classes, collected from various online sources. Our approach, as shown in figure 1, involves several steps: first, we pre-process the data to improve the classification. As the EmoHD dataset is unbalanced, we then perform resampling to balance the data. Finally, we classify emotions by experimenting with three methods: CNN, GRU, and ZSL.



Figure 1: Sentiment classification and analysis on imbalanced Health-Related textual data framework

3.1. Data preprocessing

To classify the text data in the EmoHD database, it is necessary to perform a series of pre-processing steps to remove any noise present in the data and improve the classification accuracy. The main steps included are:

- **Lowercase conversion:**
Converting all terms to lowercase is necessary to avoid duplicates, as even though the meaning of words such as "Health" and "health" are the same, they are treated as separate lexical units if they are in different cases.
- **Removal of punctuation:**
Removing all punctuations from the text that do not provide any useful information is a process that helps to improve the performance of data classification.
- **Remove stop words:** These are words that are very common in the language being studied but do not provide any informative value for understanding the "meaning" of a document or corpus, so they are removed.
- **Removal of rare and common words:**
To avoid the negative impact that rare and common words can have on the classification,

we remove them by counting their frequency of appearance in the text, this helps in reducing the noise generated by them in the text.

- **Lemmatization:**

It consists of replacing each word with its canonical form, for example, the word "known" is replaced with its canonical form "know". This step is useful for the thematic classification of texts because it treats different variants resulting from the same form or root as a single word.

- **Tokenization:**

It's the act of parsing text into tokens, in other words, the text is segmented into linguistic units such as words, punctuations, numbers, alphanumeric data... Each element corresponding to a token which will be useful for the to analyse.

3.2. Data re-sampling

The EmoHD dataset has six class labels: Angry (1343), Excited (1215), Fear (742), Happy (522), Sad (358), and Bored (22). As the data is unbalanced, it is necessary to perform resampling to balance the data. Our approach involves using naive oversampling by randomly duplicating observations from the minority class in order to increase their representation. This is done by re-sampling with replacement.

3.3. Classification models

- **Convolution Neural Network (CNN)**

Convolutional Neural Networks (CNNs) are types of deep, feedforward artificial neural networks composed of layers of nodes, including an input layer, one or more hidden layers, and an output layer [4]. Each node connects to another node and has an associated weight and threshold. If the output of an individual node exceeds the designated threshold value, that node is activated, and data is passed to the next layer of the network, otherwise, no data is passed to the next network layer. In our approach, we implemented a CNN with 5 layers, which are:

- **The Embedding Layer:** This is the input layer that maps words in the text to real-valued vectors. It takes three inputs: the vocabulary size dimension of each embedded word (input_dim), the maximum number of words in the vocabulary (max_feature), and the maximum length of a sequence (input_length).
- **Spatial Dropout Layer:** This layer is used to reduce overfitting during the training of the model by applying random deactivation at each epoch. This means that during each pass (forward propagation), the model learns with a configuration of different neurons activating and deactivating randomly.
- **1D Convolution Layer (Temporal Convolution):** This layer creates a convolution kernel that is convolved with the input over a single spatial (or temporal) dimension to produce a tensor of outputs.
- **MaxPooling1D Layer:** This layer down-samples the input representation by taking the maximum value over a spatial window of size pool_size

- **The Dense Layer:** This is the output layer, which implements the activation function. In our case, we used the softmax which is often used during the multiclass problem as in our case function.

Next, we have compile our model. Compiling the model takes three parameters: optimizer, loss and metrics. The optimizer controls the learning rate. We will be using 'adam' as our optimizer. Adam is generally a good optimizer to use for many cases. The adam optimizer adjusts the learning rate throughout training. The learning rate determines how fast the optimal weights for the model are calculated. We used 'categorical_crossentropy' for our loss function. This is the most common choice for classification. A lower score indicates that the model is performing better. And we used the 'accuracy' metric to see the accuracy score on the validation set when we train the model.

- **Gated Recurring Unit (GRU)**

GRU (Gated Recurrent Unit) is an improved version of the standard recurrent neural network that was introduced in [11]. It has some advantages over long-term memory (LSTM) in certain cases, such as using less memory and being faster.

GRUs solve the leakage gradient problem of a standard RNN by using two gates, the update gate and the reset gate. These gates are two vectors that decide what information is allowed to pass to the output and can be trained to retain information from far back in time. This allows for relevant information to be passed along a chain of events for better predictions.

The GRU model implemented in our approach has 3 layers: an embedding layer (defined in the previous section), a GRU layer (this layer is a fully connected layer with Gated recurrent units instead of simple neurons (we used 64 GRU), which are an improved version of standard recurrent neural network. The GRU is similar to a long short-term memory (LSTM) but with fewer parameters), and a dense layer with a softmax function activation.

Next, we compiled our model with the three parameters: optimizer, loss, and metric (previously quoted in the CNN model).

- **Zero Shot Learning (ZSL)**

Zero-shot learning (ZSL) is the ability to perform a task without any prior training examples. It is used to build models for classes that have not yet been labeled. ZSL transfers knowledge from known classes to new classes using class attributes as the basis of this transfer. The process of ZSL has two phases:

- **Training:** Gaining knowledge about the attributes
- **Inference:** Utilizing the knowledge to classify examples into new classes.

In our approach, we implement zero-shot classification by using the transformer library on the EmoHD dataset. Our proposed model takes the candidate labels ("Angry", "Bored", "Happy", "Sad", "Excited", "Fear") from the EmoHD dataset and the text vectorized by the sentence transformer. The output of the zero-shot model is a list of emotion labels mapped to their probabilities for the given input text.

Table 1

F1-score implemented models

Model	CNN	GRU	ZSL
with re-sampling data	0.45	0.40	0.20
without re-sampling data	0.82	0.80	0.31

4. Experiments and results

In this research, we evaluated the performance of GRU, CNN, and Zero-Shot Learning models on the EmoHD dataset. Additionally, we analyzed the labels of the EmoHD dataset using Zero-Shot Learning. For our testing procedure, we split the data into a training set (80%) and a validation set (20%).

To evaluate the performance of each model, we used the F1-score, which is a commonly used evaluation metric that is more appropriate for datasets with imbalanced classes than accuracy. The F1-score is a combination of precision and recall, and takes into account both false positives and false negatives. The formula for the F1-score is represented in equation 1.

$$F1 - score = 2 \times \frac{TP}{(2 \times TP + FP + FN)} \quad (1)$$

We trained various machine learning algorithms as described in the "Classification models" section and evaluated their performance on the EmoHD dataset on 12 epochs. In our study, we evaluated the classification models with and without re-sampling the data. The results of these evaluations are presented in Table 1.

From the results in table 1, it can be seen that the CNN model achieved the highest score of 82% while the GRU model had a score of 80% after re-sampling the data. Thus, we concluded that re-sampling improves the classification performance.

We also noticed poor performance of the Zero-shot Learning model. To understand these results, we randomly selected 50 text samples from the EmoHD dataset and classified them according to the output class labels (Angry, Excited, Happy, Fear, Sad, Bored) to compare the results of the ZSL model with the existing labels in EmoHD. An example of this can be seen in Table 2.

From the 50 samples, only 12 predictions matched the labels of EmoHD. We also observed that the confusing labels were particularly around the "Excited" label which was frequently confused with other labels.

5. Conclusion and perspectives

In this study, we proposed a method for sentiment analysis using unbalanced textual health data. Our research focused on examining the effect of unbalanced data on text classification and found that re-sampling improves classification performance. Additionally, we found that the CNN model did not perform better than other models in terms of F1-score. Furthermore, we

Table 2

Example of label comparison with ZSL.

Text	Score label ZSL	EmoHD label
karachi dengue fever claim another life frontier	0.278:(Sad)	Fear
post may fp report karachi woman died dengue	0.270:(Fear)	
fever private hospital karachi total number death	0.188:(Excited)	
fatal fever karachi reached two according detail	0.096:(Angry)	
deceased woman identified balqees bin qasim lo-	0.095:(Bored)	
cality karachi toll infected patient city reached	0.070:(Happy)	
since beginning year according dengue surveil-		
lance cell total number dengue case far reported		
karachi month may.		

discovered that the similarities in semantics among the labels used in the EmoHD dataset led to confusion, particularly with the term "Excited" which can be mistaken as both "angry" and "happy". As next steps, we plan to continue our experiments in sentiment analysis and explore the possibility of creating new textual databases from web data.

References

- [1] Z. Zhang, V. Saligrama, Zero-shot learning via semantic similarity embedding, CoRR abs/1509.04767 (2015). URL: <http://arxiv.org/abs/1509.04767>. arXiv:1509.04767.
- [2] S. Srinivasan, R. Sangwan, C. Neill, T. Zu, Twitter data for predicting election results: Insights from emotion classification, IEEE Technology and Society Magazine 38 (2019) 58–63. doi:10.1109/MTS.2019.2894472, publisher Copyright: © 2019 IEEE.
- [3] R. Marcec, R. Likic, Using twitter for sentiment analysis towards astrazeneca/oxford, pfizer/biontech and moderna covid-19 vaccines 98 (2022) 544–550. URL: <https://pmj.bmj.com/content/98/1161/544>. doi:10.1136/postgradmedj-2021-140685.
- [4] Analyse des sentiments à l'aide d'un réseau de neurones convolutifs (????) 2359–2364. doi:10.1109/CIT/IUCC/DASC/PICOM.2015.349.
- [5] N. Azam, T. Ahmad, N. U. Haq, Automatic emotion recognition in healthcare data using supervised machine learning, PeerJ Computer Science 7 (2021) e751.
- [6] G. B. Mohammad, S. Potluri, A. Kumar, R. Tiwari, R. Shrivastava, S. Kumar, K. Srihari, K. Dekeba, et al., An artificial intelligence-based reactive health care system for emotion detections, Computational Intelligence and Neuroscience 2022 (2022).
- [7] M. Mustafa, T. H. A. Soliman, A. I. Taloba, M. F. Seedik, Arabic aspect based sentiment analysis using bidirectional GRU based models, CoRR abs/2101.10539 (2021). URL: <https://arxiv.org/abs/2101.10539>. arXiv:2101.10539.
- [8] L. Shu, H. Xu, B. Liu, J. Chen, Zero-shot aspect-based sentiment analysis, CoRR abs/2202.01924 (2022). URL: <https://arxiv.org/abs/2202.01924>. arXiv:2202.01924.
- [9] Analyse des sentiments à l'aide de gru modifié, in: Actes de la quatorzième conférence internationale 2022 sur l'informatique contemporaine, ????, p. 356–361. URL: <https://doi.org/10.1145/3549206.3549270>. doi:10.1145/3549206.3549270.

- [10] S. G. Tesfagergish, J. Kapociute-Dzikiene, R. Damasvicius, Zero-shot emotion detection for semi-supervised sentiment analysis using sentence transformers and ensemble learning, *Applied Sciences* 12 (2022). URL: <https://www.mdpi.com/2076-3417/12/17/8662>. doi:10.3390/app12178662.
- [11] J. Chung, Ç. Gülçehre, K. Cho, Y. Bengio, Empirical evaluation of gated recurrent neural networks on sequence modeling, *CoRR* abs/1412.3555 (2014). URL: <http://arxiv.org/abs/1412.3555>. arXiv:1412.3555.