UNIVERSITY OF MISKOLC
FACULTY OF MECHANICAL ENGINEERING AND INFORMATICS

INVESTIGATING THE DYNAMIC CONNECTION BETWEEN LANGUAGE AND SOCIAL MEDIA IN SENTIMENT ANALYSIS

PHD THESES

Prepared by

Aadil Gani Ganie

ENGINEERING OF INFORMATION TECHNOLOGY (BSC),
ENGINEERING OF COMPUTER APPLICATIONS (MSC)

JÓZSEF HATVANY DOCTORAL SCHOOL FOR COMPUTER SCIENCE AND ENGINEERING
SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY
DECEMBER AND 2023

Head of Doctoral School

Dr. Jenő Szigeti
Full Professor

Scientific Supervisor

Dr. Samad Dadvandipour
PhD, Associate Professor

Acknowledgments

I commence by expressing my profound gratitude to the divine presence for bestowing upon me the fortitude and resilience necessary to traverse this profound academic odyssey and accomplish this PhD thesis.

My deepest appreciation extends to my parents, whose boundless love, unwavering support, and relentless encouragement have been the cornerstone of my academic achievements. Their sacrifices and unwavering dedication to my education have been pivotal in my scholarly success.

I am indebted to my supervisor (Dr. Samad Dadvandipour), whose mentorship and counsel have been invaluable throughout this research expedition. Dr. Samad's wealth of expertise, profound knowledge, and unwavering support have profoundly shaped this thesis, and for that, I am profoundly grateful.

I want to express my gratitude to my colleagues, whose camaraderie and support have been a wellspring of inspiration. Their enduring friendship has imbued this academic journey with profound meaning and camaraderie.

I acknowledge with deep respect the steadfast dedication of the faculty members and reviewers. Their constructive critiques and sagacious guidance have played a pivotal role in refining my work, and for that, I am sincerely appreciative of their commitment to scholarly excellence.

Lastly, my heartfelt gratitude goes out to all those who, in myriad ways, supported me during the course of this PhD thesis. Their encouragement and support have been immeasurable, and I recognize that my journey has been enriched by their presence.

I am profoundly thankful for the unwavering support and guidance I have received throughout this odyssey. I earnestly hope that this PhD thesis stands as a testament to the collective diligence and dedication of all those who have contributed to its fruition.

Table of Contents

Preface

This dissertation aims to explore the intricate interaction between language and social media within the domain of sentiment analysis. The exponential growth of social media data has necessitated the development of sophisticated techniques for automated sentiment analysis of textual content. This research endeavors to address five critical research questions that delve into various aspects of sentiment analysis of social media text.

The first inquiry investigates the efficacy of fine-tuning pre-trained transformer models on social media text for sentiment analysis, as compared to training models from scratch on social media data. Fine-tuning of pre-trained models has gained significant traction in the natural language processing arena, and optimizing such models for specific tasks can enhance their performance. However, the compatibility of these models with sentiment analysis of social media text remains unclear. Hence, this research seeks to provide a nuanced understanding of the strengths and limitations of fine-tuning pre-trained models on social media data. The study also explores the integration of clustering attention mechanisms with autoencoders in Natural Language Processing (NLP). The study achieves promising adaptability and performance across domains, demonstrating enhanced accuracy in financial sentiment analysis, fake news detection, and hate speech detection, with a focus on efficiency in resource-constrained environments.

The second research question focuses on the impact of informal language like emoticons, hashtags (Twitter), and slang- on the accuracy of sentiment analysis models applied to social media text. Informal language is ubiquitous in social media text, presenting a significant challenge for traditional sentiment analysis models designed for more formal text. This research aims to identify the role of informal language in sentiment analysis of social media text and identify methods to enhance the performance of sentiment analysis models in the presence of such language.

The third inquiry investigates how filter size and the number of filters affect the performance of Convolutional Neural Network (CNN)-based sentiment analysis models on three datasets: Amazon review, Amazon food review, and hate speech detection. The findings reveal that the optimal combination of filter size and number of filters varies with the dataset employed. Fine-tuning these parameters can improve the accuracy of the CNN model for sentiment analysis. The study underscores the criticality of considering these factors to attain optimal performance in sentiment analysis.

The fourth research question explores the use of social media posts to get valuable insights into the mental health of individuals and communities, using machine learning models capable of accurately predicting mental health status from social media text. With mental health being a growing concern, social media posts hold immense potential in providing a window into the mental health status of individuals and communities.

The fifth research question delves into the relationship between hate speech and mental health using logistic regression models trained on two preprocessed datasets - mental health and hate speech. The visualization of hate speech distribution in the mental health dataset suggests a potential relationship between hate speech and mental health, particularly in individuals with anxiety and bipolar disorders.

The final research question proposes a novel approach to gradient descent called "Sample Gradient Descent," which entails choosing an adequate data sample and subjecting it to batch gradient descent. The study uses Principle Component Analysis (PCA) to select the sample, resulting in faster convergence rates and reduced computation times compared to conventional batch gradient descent. The proposed method offers potential utility in various domains, including machine learning and optimization problems. What impact does fine-tuning have on the performance and interpretability of large language models across diverse textual domains, and how does it contribute to enhancing their adaptability to specific tasks. Overall, this dissertation offers a comprehensive analysis of sentiment analysis in social media and provides valuable insights into future research directions.

1. Natural Language Processing

Natural Language Processing (NLP) is a branch of artificial intelligence and computer science that studies natural language interactions between computers and people. It entails creating models and algorithms to comprehend, evaluate, produce, and work with human language. Natural language processing (NLP) finds wide-ranging uses in domains like question-answering systems, machine translation, sentiment analysis, text classification, and text summarization. A branch of linguistics, computer science, and artificial intelligence called natural language processing (NLP) studies how computers and human (natural) languages interact. Ultimately, it aims to make computers capable of reading, understanding, and producing text that is human-like.

The scientific investigation of language and its structural components, including grammar, semantics, and phonetics, is referred to as linguistics. Classical linguistics involves the creation and examination of rules governing language use. While significant progress has been made in the development of formal methods for syntax and semantics, the complexities of natural language comprehension often defy precise mathematical formalization. A linguist is generally understood to be a person who studies language, although the term may also be applied to individuals who engage in fieldwork and empirical research in the field. The use of mathematical techniques and principles to analyze and understand natural language is sometimes referred to as mathematical linguistics and may involve the application of discrete mathematical formalisms and theories, such as formal languages and automata theory, to the study of language. Computational linguistics, as a subfield of linguistics and computer science, involves the application of computational and statistical methods to the study of natural language and its structural and functional properties. With the advent of big data and advanced computing technologies, computational linguists have been able to leverage large datasets of natural language text and speech to discover novel patterns and insights that were previously beyond the reach of traditional rule-based approaches. In recent years, the statistical approach to natural language processing, which relies on probabilistic models and machine learning algorithms, has emerged as the dominant paradigm in the field, supplanting earlier rule-based approaches that were limited by their computational complexity and inflexibility. This shift towards a more empirical and data-driven approach has led to the use of the term "natural language processing" (NLP) to describe the field of computational linguistics, which seeks to develop computational models and algorithms that can facilitate the automatic analysis, interpretation, and generation of natural language data.

2. Sentiment Analysis

Sentiment analysis is a branch of natural language processing (NLP) that focuses on automatically recognizing and extracting subjective data from text.

Determining whether a sentiment is positive, negative, or neutral in a textual document—like a news article, social media post, or product review—is the aim of sentiment analysis. Numerous applications, including market, political, and customer sentiment analysis, can benefit from this information.

Sentiment analysis, commonly referred to as opinion mining, is a branch of natural language processing (NLP) that uses computational methods to locate and extract subjective data from textual data in order to quantify and characterize the opinions, feelings, and attitudes that are expressed in the text. (Pang & Lee, 2008). This may involve the use of supervised or unsupervised machine learning algorithms, lexical resources, and computational linguistic methods to analyze the syntactic, semantic, and pragmatic features of the text, as well as contextual and cultural factors that may influence the interpretation of the text (B. Liu, 2012). At the most basic level, sentiment analysis involves the classification of text as positive, negative, or neutral, based on the presence or absence of certain keywords or phrases that are indicative of positive or negative sentiment (Hu & Liu, 2004) More advanced approaches may involve the identification and classification of specific emotions or affective states, such as joy, sadness, anger, fear, or surprise (Camras & Plutchik, 1980), or the quantification of the overall valence or intensity of the sentiment expressed in the text (Baccianella et al., 2008).

At the most basic level, sentiment analysis involves the classification of text as positive, negative, or neutral, based on the presence or absence of certain keywords or phrases that are indicative of positive or negative sentiment. More advanced approaches may involve the identification and classification of specific emotions or affective states, such as joy, sadness, anger, fear, or surprise, or the quantification of the overall valence or intensity of the sentiment expressed in the text. In addition to its applications in social media analysis, market research, and customer service, sentiment analysis has also been used in various domains, including politics, journalism, and psychology, to gain insight into the attitudes and opinions of large groups of people.

2.1 Social Media and Sentiment Analysis

The exponential growth in digital text data and machine learning advancements have driven the field of Natural Language Processing (NLP) to rapid growth in recent years. One of the key applications of NLP is sentiment analysis, which aims to determine the polarity of a given piece of text, whether it is positive, negative, or neutral. Sentiment analysis has wide-ranging applications, including opinion mining, marketing research, customer relationship management, and monitoring public opinion, among others. Social media has emerged as one of the most important sources of user-generated content and is an especially rich source of information for sentiment analysis. Users can freely and publicly express their opinions and feelings on social media platforms like Facebook, Instagram, and Twitter, which makes them a rich and varied source of data for sentiment analysis.

The growth of social media and its impact on society have made sentiment analysis of social media data a critical research area. The challenge of sentiment analysis of social media data lies in the informal nature of the language used in such texts. Social media users often use informal language, including emoticons, hashtags, slang, and misspelled words, which can affect the accuracy of sentiment analysis models. Additionally, the multilingual and multicultural nature of social media can present challenges for sentiment analysis models, as sentiment expressions can vary greatly across different languages and cultures.

To address these challenges, researchers have explored different approaches to sentiment analysis of social media data. One approach is to fine-tune pre-trained transformer models, such as BERT and RoBERTa, on social media data. Another approach is to train models from scratch on social media data, leveraging the unique characteristics of the language used in social media texts. The performance of these models is then compared to determine which approach is more effective.

3. Research Questions

The aim of this thesis is to explore the intersection of language and social media in sentiment analysis. The following research questions are addressed:

- Can pre-trained transformer models be fine-tuned effectively on social media text for sentiment analysis, and if so, how does this compare to training models from scratch on social media data?
- How can clustering attention models, enriched by autoencoder-based dimensionality reduction, be optimized for dynamic clustering techniques and scaled to larger datasets while maintaining interpretability in NLP applications?
- How does the presence of informal language, such as emoticons, hashtags, and slang, impact the performance of sentiment analysis models on social media text?
- What is the significance of filter size and number of filters in Convolutional Neural Networks for sentiment analysis.
- Investigating the Use of Machine Learning for Assessing Mental Health through Analysis of Social Media Posts.
- Is there a relationship between hate speech and mental health, and if so, what are the potential mechanisms underlying this relationship?

- How does "Sample Gradient Descent" compare to conventional batch gradient descent in terms of convergence rates and computation times, and what are the potential applications of this method in machine learning and optimization problems?

- What impact does fine-tuning have on the performance and interpretability of large language models across diverse textual domains, and how does it contribute to enhancing their adaptability to specific tasks?

- How do different hyperparameters, such as learning rate, batch size, and layer-specific tuning, influence the fine-tuning process of large language models, and what optimal configurations exist for various natural language processing tasks

4. Pre-trained Transformer Models for Sentiment Analysis on Social Media Text

One kind of deep learning architecture that has gained popularity for natural language processing tasks, such as sentiment analysis, is pre-trained transformer models. Compared to training from scratch, these models can be refined to accomplish particular tasks using a far smaller amount of data because they have already been pre-trained on enormous volumes of text data. The Transformer architecture, which was presented in the 2017 paper "Attention Is All You Need" by Vaswani et al., serves as the foundation for Transformer models. (Vaswani et al., 2017). The key innovation of the Transformer architecture is the attention mechanism, which allows the model to attend to different parts of the input sequence when making predictions. This allows the model to effectively capture long-range dependencies in the input data, making it well-suited for tasks such as language translation.

Fine-tuning pre-trained transformer models for sentiment analysis on social media text is an effective approach that leverages the powerful language representation capabilities of these models to solve NLP tasks. The main idea behind fine-tuning is to use the pre-trained weights of a large transformer model as the starting point for training on a new task, rather than training the model from scratch. This approach has been widely adopted in NLP because it can significantly reduce the amount of data and computation required to train a model, while still allowing the model to perform well on new tasks. Sentiment analysis is a popular NLP task that involves classifying the sentiment of a text as positive, negative, or neutral. In the context of social media text, sentiment analysis can be used to identify patterns and trends in public opinion, understand the tone and mood of social media conversations, and inform decision-making in various industries. To fine-tune a pre-trained transformer model for sentiment analysis, the first step is to prepare a training dataset that consists of social media text and corresponding sentiment labels. This dataset can be annotated manually or obtained from existing annotated datasets.

The identification of online hate speech is a critical issue for the field of natural language processing (NLP), particularly as social media has amplified this phenomenon by providing a virtual platform for online harassment. This study aims to address this issue by utilizing the trolling aggression and cyber-bullying dataset from the shared tasks workshop. To accomplish this, the study employs an extreme pre-processing methodology and implements an ensemble approach in the model building process. The study also evaluates the performance of existing algorithms such as random forest, logistic regression, and multinomial Naïve Bayes. The results demonstrate that logistic regression is the most efficient algorithm, achieving an accuracy of 57.91%. However, the ensemble bidirectional encoder representation from transformers demonstrates even more promising results, yielding a precision of 62%, outperforming most existing models. For the purpose of this study, the data was obtained from the shared task on aggression identification dataset, which was organized at the "Trolling, Aggression, and Cyber-Bullying Workshop" (Ganie & Dadvandipour, 2021). The training data consisted of 10,799 Facebook comments that were randomly selected and annotated into three categories: Overly Aggressive (OAG), Covertly Aggressive (COA), and Non-Aggressive (NAG). The validation or test data comprised of 1200 samples.

5. Enhancing Language Comprehension: An Innovative Approach Incorporating Clustered Attention in Transformer Models

This study pushes the boundaries of text classification through the introduction of an innovative clustering-based attention mechanism seamlessly integrated with embeddings, autoencoders, and inspired by transformer architectures. Evaluating its impact on financial sentiment analysis, fake news detection, and hate speech detection datasets, our research reveals consistent improvements in accuracy compared to conventional attention mechanisms. The incorporation of autoencoders contributes to performance refinement by reducing dimensions and optimizing both time and space efficiency. Notably, our model maintains comparable performance while exhibiting significantly reduced complexity, as evident from execution times and parameter counts. This research makes a significant contribution to the evolving landscape of attention mechanisms in natural language processing, showcasing the synergistic benefits of clustering strategies, embeddings, autoencoders, and transformer-inspired architectures in elevating the effectiveness of text classification models. The amalgamation of these components emerges as a promising paradigm, paving the way for advancements in precision, interpretability, and efficiency within the realm of text classification.

In our quest to enhance attention mechanisms through clustering, we introduce a novel objective function that refines the conventional attention scoring mechanism.

Our focus is on the key (K) and query (Q) matrices, integral components in attention computation. Let's start with the mathematical background. Consider a set of data points D with dimensions D. Each data pointx_i can be represented as a vector x_i in R^D. In the attention mechanism, the attention score α_{ij}between a query $q_i$and key $k_j$is traditionally computed as the scaled dot product

$$\alpha_{ij} = \frac{q_i . k_j}{\sqrt{D}}$$

This formulation ensures that the gradients remain stable during training. To incorporate clustering into the attention mechanism, we propose clustering the matrices $K$and Q. Let $K^c$and Q^cdenote the clustered version of $K$and . The clustering operation can be expressed as:

$$K^c = Cluster(K)$$

$$Q^c = Cluster(Q)$$

This operation groups similar vectors together, fostering a more structured attention computation. Now, we define our novel objective function as the mean squared error between the attention scores with and without clustering:

$$l = \frac{1}{n} \sum_{i=1}^{N} \sum_{j=1}^{N} \left(\alpha_{ij}(K^c, Q^c) - \alpha_{ij}(K, C) \right)^2$$

where N is the number of data points.

Algorithm: Custom Attention Training Loop

1. Input:

1. **num_epochs:** Number of training epochs.
2. **data_batches:** Batches of input data for each iteration.

2. Initialize Model Parameters:

1. Initialize or load the model parameters.

3. Define Functions:

1. **scaled_dot_product_attention(q, k, v):** Define the scaled dot-product attention mechanism.
2. **cluster_matrix(matrix):** Define a function to perform clustering on the input matrix.
3. **custom_attention(q, k, v):** Define the custom attention mechanism using clustered matrices and compute the loss as the mean squared error (MSE) between the attention mechanisms.

4. Training Loop:

A. For each ***epoch*** in the range ***num_epochs***:
   . For each ***batch*** in ***data_batches:***
      a. Preprocess the batch to obtain query ***(q)***, key ***(k)***, and value ***(v)*** matrices.
      b. Compute the clustered attention and loss using the ***custom_attention*** function
      c. Update the model parameters using the loss gradient (not explicitly shown in the code snippet).

5. Output:

- The trained model parameters.

Conclusions

This study was conducted with the primary objective of delving into the prevalence of online harassment within the digital media landscape, employing the TRAC-1 dataset sourced from the trolling, aggression, and cyber-bullying workshop's shared task. Distinguishing itself from prior research endeavors, which often relied on partially preprocessed data, this study implemented an extensive preprocessing procedure. This encompassed advanced techniques such as contractions handling, lemmatization, stemming, and stop-word removal to ensure a robust foundation for subsequent analyses. The initial application of traditional machine learning algorithms, including logistic regression and Naïve Bayes, yielded unsatisfactory accuracy rates. Recognizing the need for performance enhancement, the study pivoted towards an ensemble approach featuring fine-tuned pre-trained BERT models. Impressively, this approach outperformed existing state-of-the-art models, underscoring the significance of leveraging advanced language models for online harassment detection. Nonetheless, the study brought to light the nuanced challenges associated with hyperparameter tuning, as an observed increase in batch size and learning rate led to overfitting and a subsequent decline in accuracy.

These findings not only contribute to the current understanding of online harassment detection but also pave the way for future investigations. The identified avenues for research emphasize the importance of optimizing model performance through meticulous hyperparameter tuning. This involves a delicate balance between batch size and learning rate, a nuanced task that warrants further exploration. Moreover, the study delved into the integration of clustering attention mechanisms, complemented by the incorporation of autoencoders, within the broader domain of Natural Language Processing (NLP). The adaptability and performance demonstrated across diverse domains, including financial sentiment analysis, fake news detection, and hate speech detection, underscore the versatility of the clustering attention model. The model's efficiency in both time and space complexity, especially when enriched by autoencoder-based dimensionality reduction, positions it as a compelling choice for resource-constrained environments.

The study's emphasis on scalability, interpretability, and further experimentation with dynamic clustering techniques reveals the depth of potential avenues for future exploration. The insights gleaned from this work not only advance the understanding of attention mechanisms in NLP but also present an intriguing intersection between deep learning, autoencoders, and language processing, inviting further investigations into the development of more robust and interpretable models for real-world applications.

6. The Impact of Informal Language on Sentiment Analysis Models on Social Media Text

In social media, informal language has become increasingly prevalent in online communication. It is characterized by the use of non-standard grammatical structures, informal vocabulary, and unconventional spellings and punctuation. Informal language in social media text is important to study because it provides insights into the ways in which users interact with one another, as well as the cultural, social, and linguistic norms that are emerging in this new digital landscape. Research has shown that the use of informal language in social media text is influenced by several factors, including age, gender, education level, and cultural background. For example, a study by (Hill, 2016) found that younger users and users with lower levels of education were more likely to use informal language in their online interactions, while users from different cultural backgrounds showed differences in their use of linguistic elements such as emoticons and abbreviations. In addition to exploring the usage patterns of informal language in social media text, researchers have also investigated the impact of informal language on the meaning and interpretation of messages. For instance, studies have shown that the use of informal language can influence the perceived level of politeness, credibility, and intentionality in online communication (Herring et al., 2013). It is critical to keep researching the use and effects of informal language in social media text because of the growing significance of social media in our daily lives and the growing prevalence of informal language in online interactions.

7. The Effect of Informal Language on Sentiment Analysis Models

This investigation aimed to explore the effect of informal language on the efficacy of sentiment analysis models applied to social media text. The research utilized a Convolutional Neural Network (CNN) approach, and the model was developed and trained on three different datasets: a sarcasm corpus, a sentiment corpus, and an emoticon corpus. The experimental design held the model architecture constant and trained the model on 80% of the data, then evaluated its performance on the remaining 20%. The outcomes indicated that the model showed an accuracy of 96.47% on the sarcasm corpus, with the lowest accuracy for class 1. The sentiment corpus elicited an accuracy of 95.28% from the model. The accuracy increased marginally to 95.1% with the integration of the sentiment and sarcasm datasets, and by slightly more than that with the addition of the emoticon corpus, to 95.37%. These findings imply that informal language has little effect on how well sentiment analysis models perform when applied to text from social media, however the addition of emoticon data can somewhat increase accuracy.

Conclusion

In the pursuit of comprehending the impact of informal language, encompassing elements like emoticons and slang, on the effectiveness of sentiment analysis models applied to social media text, this study undertook a meticulous examination. The outcomes of our investigation reveal a noteworthy accomplishment, with the model achieving a commendable accuracy rate of 96.47% when applied to the sarcasm dataset. However, it is noteworthy to mention that the model exhibited its weakest performance in class 1 within the sarcasm dataset. When the model was exclusively tested on the sentiment dataset, it demonstrated a robust accuracy of 95.28%. Intriguingly, the amalgamation of sarcasm and sentiment data yielded an overall accuracy of 95.1%, suggesting a cohesive performance across diverse linguistic dimensions. Furthermore, the augmentation of our model with emoticon data resulted in a marginal improvement, pushing the accuracy to 95.37%. These findings collectively suggest that the influence of informal language on sentiment analysis model performance, particularly in the realm of social media text, appears to be somewhat limited. However, the subtle enhancement observed with the inclusion of emoticon data emphasizes the nuanced nature of language in the digital landscape.

While these results provide valuable insights, the horizon of potential research endeavors remains expansive. Future investigations could delve into a more granular exploration of the impact of diverse forms of informal language, including but not limited to emojis and hashtags, on sentiment analysis models. The incorporation of these elements may unearth additional layers of complexity in the linguistic landscape of social media text. Moreover, the exploration of alternative model architectures and machine learning methodologies stands as a promising avenue for further refinement. Considering the dynamic nature of social media language, the utilization of advanced techniques such as recurrent neural networks and transformer networks may unlock new dimensions of accuracy and adaptability in sentiment analysis models.

In essence, while our study sheds light on the nuanced relationship between informal language and sentiment analysis model performance, it is clear that there is more to unravel. By expanding the scope of investigation to encompass a broader array of informal language elements and employing cutting-edge modeling techniques, future studies hold the potential to contribute significantly to the evolving landscape of sentiment analysis in the dynamic and ever-evolving realm of social media text.

8. Optimizing Filter Size and Number of Filters for Sentiment Analysis using CNN: A Comparative Study across Different Datasets

This chapter examines how the number and size of filters affect a Convolutional Neural Network (CNN) based sentiment analysis model's performance. The study was conducted on three datasets: Amazon review, Amazon food review, and hate speech detection. The results of the experimentation established that filter size and number of filters have a substantial impact on the performance of a CNN-based sentiment analysis model. The results showed that the performance of the model was contingent on the dataset employed and the combination of filter size and number of filters. The Amazon food review dataset exhibited the highest accuracy at 90.49%, while the hate speech detection dataset performed relatively less with an accuracy of 88.23%. The study suggests that fine-tuning the filter size and number of filters can enhance the performance of a CNN-based sentiment analysis model. This chapter gives a thorough overview of how the number and size of filters affect sentiment analysis effectiveness with a CNN model and emphasizes how crucial it is to adjust these parameters for best results.

In our comprehensive study, we conducted a meticulous examination of the intricate interplay between filter size and the number of filters in Convolutional Neural Networks (CNNs) applied to sentiment analysis, scrutinizing three distinct datasets. The findings unveiled a discernible trend, indicating that larger filter sizes coupled with a greater number of filters generally resulted in heightened performance across the datasets, with a notable emphasis on the 4:128 ratio. However, the absence of a singular, optimal configuration underscored the nuanced balance required between filter size and the number of filters, intricacies that emerged as dataset-specific and task-dependent.

These insights, serving as invaluable guideposts, offer practitioners and researchers navigating the domain of CNNs in Natural Language Processing tasks a roadmap for optimizing their choices of filter size and quantity to attain superior performance. Yet, the pursuit of the optimal configuration remains an ongoing challenge, prompting the need for continuous research to validate and extend these findings to diverse datasets and a spectrum of Natural Language Processing applications.

Despite the strides made in understanding the relationship between filter size, filter quantity, and sentiment analysis performance, this quest for optimization delves into a complex landscape. The dataset-specific and task-dependent nature of these intricacies necessitates a nuanced approach in tailoring CNN architectures for optimal outcomes in sentiment analysis tasks. The insights derived from our study, therefore, contribute not only to the current understanding of sentiment analysis models but also pave the way for future endeavors seeking to unravel the complexities inherent in language nuances and model architecture intricacies. Particularly in the dynamic landscape of Natural Language Processing, where the challenges of social media text present unique hurdles, these insights offer a robust foundation for researchers to navigate and explore further refinements in pursuit of enhanced model performance and interpretability.

9. Investigating the Relationship between Hate Speech and Mental Health Using Natural Language Processing Techniques

This study aims to investigate the relationship between hate speech and mental health by training separate models for each dataset. The mental health dataset consisted of text data with three classes - Anxiety, Bipolar, and Normal - while the hate speech dataset had two classes - Hate and Normal. Both datasets were preprocessed and encoded before being used to train logistic regression models. The hate speech model achieved an accuracy score of 0.946, while the mental health model achieved an accuracy score of 0.930. The visualization of hate speech distribution in the mental health dataset showed a significant proportion of samples with no hate speech, particularly in the Normal class. However, the proportion of hate speech samples was relatively high in the Anxiety and Bipolar classes. The results suggest a potential relationship between hate speech and mental health, particularly in individuals with anxiety and bipolar disorders. The use of hate speech in online platforms could contribute to exacerbating the symptoms of these disorders. Further research is necessary to investigate the specific mechanisms underlying this relationship and to explore potential interventions to mitigate the negative impact of hate speech on mental health.

In summary, the proliferation of hate speech has emerged as a widespread and concerning issue in societies globally, exerting detrimental effects on mental health. This concern is particularly pronounced in online spaces, where the rapid dissemination of hate speech to a broad audience exacerbates its impact. Recent scholarly inquiries have delved into the intricate relationship between hate speech and mental health, with some studies focusing on specific marginalized groups and others taking a broader perspective to understand the overall mental health outcomes. This research paper adds to this burgeoning field by leveraging Natural Language Processing (NLP) techniques to autonomously classify instances of hate speech within a comprehensive dataset of online discourse, shedding light on their connection to mental health outcomes.

The investigation discerned a noteworthy association between exposure to hate speech and adverse mental health outcomes, especially affecting individuals from marginalized groups. These findings underscore the urgency for effective interventions aimed at mitigating the deleterious impact of hate speech on mental well-being. The incorporation of NLP and machine learning techniques for the automated detection and classification of hate speech introduces promising avenues for the development of such interventions. However, the efficacy of these approaches hinges on several critical factors, including the quality and representativeness of the training data, the ability to capture the nuanced nature of hate speech, and the delicate balance required to address concerns related to free speech and censorship. In essence, the study not only illuminates the negative consequences of hate speech on mental health but also emphasizes the imperative for interventions to alleviate this impact.

While the application of NLP and machine learning holds potential as a solution to combat hate speech, further research is indispensable to comprehensively understand the effectiveness and limitations of these approaches. The complexities surrounding the nuanced nature of hate speech necessitate ongoing inquiry to refine interventions and strike an equitable balance between safeguarding mental health and upholding principles of free expression. In conclusion, this study contributes to a deeper understanding of the intricate dynamics between hate speech and mental health, laying the groundwork for future research endeavors to refine and expand upon potential solutions in this critical area.

10. Social Media Posts as a Window into Mental Health: A Machine Learning Approach

Mental health stands as a pivotal determinant of human well-being, drawing substantial attention in recent years due to the widespread occurrence of mental health disorders and their adverse impact on individuals and society at large. Addressing this pressing concern, researchers have explored the vast reservoir of data available on social media platforms as a potential means to predict and classify mental health statuses. In our study, we meticulously analyzed three distinct datasets: the first encompassed seven classes (depression, anxiety, autism (Autism, or Autism Spectrum Disorder (ASD), is a developmental disorder characterized by challenges in social interaction, communication difficulties, and repetitive behaviors. It is a spectrum condition, meaning that individuals with autism may exhibit a wide range of symptoms and abilities.), mental health, schizophrenia (Schizophrenia is a severe mental disorder characterized by disruptions in thought processes, emotions, and perception of reality. Common symptoms include hallucinations, delusions, disorganized thinking, and impaired social or occupational functioning.), BPD

(Borderline Personality Disorder is a mental health condition marked by patterns of unstable relationships, self-image, and emotions. Individuals with BPD may experience intense mood swings, impulsivity, fear of abandonment, and difficulties with interpersonal relationships.), and bipolar (Bipolar Disorder, formerly known as manic-depressive illness, is a mood disorder characterized by extreme mood swings between periods of mania (elevated mood, increased energy) and depression (low mood, decreased energy). These mood episodes can vary in duration and intensity.)), the second featured two classes (positive and negative), and the third comprised two classes (suicide and non-suicide). Additionally, a final dataset was curated, comprising 14 classes, with seven falling under the non-suicidal subset and seven within the suicidal subset. Our approach involved the implementation of logistic regression, support vector machines, and multinomial naive Bayes for classification and prediction. The performance of our models was rigorously evaluated using receiver operating characteristic (ROC) curves and confusion matrices, with the logistic regression model outshining others, achieving an impressive accuracy rate of 80%. These models were seamlessly deployed via Streamlit, offering a user-friendly interface for predicting mental health status and assessing the risk of suicidal ideation. Notably, if a social media post falls within the category of the suicidal subset class, an intelligent chatbot (GPT2) is activated, aiming to engage individuals exhibiting suicidal ideation and mitigate the likelihood of suicide. Our research not only serves as a valuable tool for mental health professionals but also holds the potential for extension to other platforms, effectively addressing the urgent imperative to detect and intervene in mental health issues and suicidal ideation.

Concluding our in-depth investigation, we emphasize the transformative potential residing within advanced machine learning methodologies when applied to forecast individuals' mental health conditions based on a meticulous analysis of their social media activity. This comprehensive study not only validates the practicality of utilizing social media data for predictive analytics concerning mental health but also extends its scope to encompass a thorough assessment of the associated risk of suicidal ideation.

The purposeful application of cosine similarity, a pivotal component of our methodological framework, serves a dual function: introducing additional classes to enrich the dataset and standing out as a crucial contributor to the substantial enhancement observed in the overall performance of our machine learning model. This strategic integration positions our methodology as a robust and sophisticated tool in the domain of mental health prediction. Furthermore, the integration of a GPT-2 chatbot introduces an innovative dimension to our study, manifesting promise as a dynamic engagement tool, particularly for individuals grappling with suicidal thoughts. This augmentation opens avenues for potential risk mitigation, highlighting the broader societal implications of our research beyond conventional predictive accuracy metrics.

Empirical evidence from our rigorous experiments showcases the superior predictive prowess of logistic regression over alternative algorithms, surpassing support vector machines and multinomial naive Bayes with a noteworthy accuracy rate of 80%. This finding not only substantiates the effectiveness of our chosen approach but also prompts further exploration into the intricacies of algorithmic selection in the context of mental health prediction. The deployment of our predictive models through the user-friendly Streamlit interface marks a significant stride toward practical accessibility. This ensures that mental health professionals and stakeholders can seamlessly integrate our methodology into their practices, fostering a proactive approach to mental health status prediction and risk assessment for suicidal ideation. As a foundational contribution to the field, our research not only stands as a valuable resource for mental health professionals but also holds promise for broader applications across diverse platforms. Addressing the critical imperative of early identification and intervention in mental health concerns and suicidal ideation, our work lays the groundwork for future investigations. Prospective research avenues could delve into exploring alternative machine learning algorithms, integrating additional data sources, leveraging transfer learning techniques, and conducting a thorough examination of the ethical and societal implications associated with the use of social media data for mental health prediction.

In summation, our multifaceted approach stands as a beacon at the intersection of cutting-edge technology and the imperative societal need for mental health support. It is poised to make a substantial and lasting impact, contributing significantly to the ongoing discourse surrounding the understanding, prevention, and mitigation of mental health challenges in contemporary society.

11. From Big Data to Smart Data: A Sample Gradient Descent Approach for Machine Learning and Fine Tuning the LLM's

This research presents an innovative approach to gradient descent known as "Sample Gradient Descent". This method is a modification of the conventional batch gradient descent algorithm, which is often associated with space and time complexity issues. The proposed approach involves the selection of a representative sample of data, which is subsequently subjected to batch gradient descent. The selection of this sample is a crucial task, as it must accurately represent the entire dataset. To achieve this, the study employs the use of Principle Component Analysis (PCA), which is applied to the training data, with a condition that only those rows and columns of data that explain 90% of the overall variance are retained. This approach results in a convex loss function, where a global minimum can be readily attained. The outcomes of our study reveal that the suggested approach exhibits accelerated convergence rates and diminished computation durations in contrast to the traditional batch gradient descent algorithm. These results underline the potential applicability of the "Sample Gradient Descent" method across diverse domains, encompassing fields from machine learning to optimization problems.

Also, we conducted a research on fine-tuning large language models (LLMs) such as phi-1.5 and Falcon series. Utilizing the AgentInstruct dataset, phi-1.5 demonstrated superior performance, achieving substantial reductions in training and evaluation losses. Falcon-7B and Falcon-RW-1B, specialized for mental health advice, exhibited nuanced training dynamics, showcasing adaptability even with limited, high-quality data. Our analysis focused on learning rate optimization, revealing its pivotal role in model convergence. Concurrently, we investigated the computational demands, emphasizing the necessity of robust infrastructure for LLM training. These findings provide crucial insights into the intricate aspects of LLM fine-tuning, offering valuable guidance for future research endeavors in the realm of large language model optimization.

11.1 Implementation

```
class GDRegressor:

    def __init__(self,learning_rate,epochs):
        self.m = 100
        self.b = -120
        self.lr = learning_rate
        self.epochs = epochs

    def fit(self,X,y):
        # calcualte the b using GD
        for i in range(self.epochs):
            loss_slope_b = -2 * np.sum(y - self.m*X.ravel() - self.b)
            loss_slope_m = -2 * np.sum((y - self.m*X.ravel() - self.b)*X.ravel())

            self.b = self.b - (self.lr * loss_slope_b)
            self.m = self.m - (self.lr * loss_slope_m)
        print(self.m,self.b)

    def predict(self,X):
        return self.m * X + self.b
```

__init__(self,learning_rate,epochs):

Initializes the object with two arguments: learning_rate and epochs.
Sets the initial values of the weights m and b.
Stores the learning_rate and epochs as object variables for later use.

fit(self,X,y):

Implements the gradient descent algorithm for a specified number of epochs.

Calculates the gradients of the loss with respect to the weights m and b.
Updates the weights m and b based on the calculated gradients and the learning rate.
Prints the final optimized values of m and b.

predict(self,X):

Predicts the output values for the given input X using the learned weights m and b.
Returns the predicted output values.

To get the sample data from the training data which explains 90% variance of data we use the following logic:

```
from sklearn.decomposition import PCA

# create a PCA object with n_components set to None to keep all components
pca = PCA(n_components=None)

# fit the PCA model to the training data
pca.fit(X_train)

# calculate the cumulative sum of explained variance ratios
cumulative_variances = np.cumsum(pca.explained_variance_ratio_)

# get the index of the first component that explains 90% of the variance
n_components = np.argmax(cumulative_variances >= 0.90) + 1

# create a new PCA object with the optimal number of components
pca = PCA(n_components=n_components)

# fit the new PCA model to the training data
pca.fit(X_train)

# transform the data to the new reduced dimensionality
X_filtered = pca.transform(X_train)
X_filtered = pd.DataFrame(X_filtered)
X_filtered
```

In conclusion, our study introduces a pioneering approach known as "Sample Gradient Descent," which harnesses the computational power of Principal Component Analysis (PCA) to strategically select a representative data sample, subsequently implementing batch gradient descent. This unique methodology not only significantly enhances computational efficiency but also maintains a commendable level of performance when compared to traditional batch gradient descent methodologies. The strategic incorporation of hyperparameters, complemented by their meticulous optimization through established techniques such as grid search or random search, contributes to the heightened adaptability and versatility of the proposed approach.

Moreover, our empirical findings shed light on the remarkable observation that the loss function of the sampled data converges at an accelerated rate, owing to the inherent convex nature of the selected data subset. This rapid convergence, in conjunction with a thorough comparative analysis against conventional gradient descent techniques, serves as a compelling testament to the effectiveness of our novel approach. We posit that the innovative "Sample Gradient Descent" technique holds immense promise for making a substantial impact across a spectrum of domains, including but not limited to machine learning and optimization. Its ability to bolster computational efficiency without compromising on performance is a notable stride forward in the realm of gradient-based optimization algorithms.

Additionally, the integration of hyperparameter tuning techniques further amplifies the adaptability and robustness of our proposed method. The meticulous experimentation conducted in this study not only enhances our understanding of large-scale language models but also unveils their inherent adaptability and efficacy across diverse tasks and datasets. These nuanced insights underscore the significant potential and practical applicability of such models in real-world scenarios, showcasing their resilience and versatility as indispensable tools in the field of data science and artificial intelligence.

12. Summary of Findings

The conclusion of this PhD thesis outlines the results of four different studies conducted on various topics related to the field of Natural Language Processing (NLP).

Study 1 investigated the prevalence of online harassment in digital media and found that an ensemble approach of fine-tuned pre-trained BERT models outperformed state-of-the-art models in terms of accuracy, but overfitting was observed with an increase in batch size and learning rate. The study demonstrates the efficacy of clustering attention mechanisms enriched by autoencoders in NLP, with superior efficiency in time and space complexity. Future research should refine the model, explore dynamic clustering techniques, and assess scalability and interpretability for real-world applications, presenting opportunities for advancing attention mechanisms in the intersection of deep learning, autoencoders, and NLP.

Study 2 and 3 The research explored the impact of informal language, including emoticons and slang, on sentiment analysis models applied to social media text. The study revealed that integrating emoticon and slang data had a marginal influence on model accuracy. Additionally, the investigation into filter size and the number of filters in CNNs emphasized the need for a delicate balance, with larger filters generally enhancing performance. These nuanced insights provide valuable guidance for NLP practitioners.

Study 4 and 5 Addressing the connection between hate speech and mental health, this study employed NLP techniques to detect hate speech in online discourse. Exposure to hate speech was linked to negative mental health outcomes, particularly affecting vulnerable populations. Advanced machine learning techniques, including the use of cosine similarity and a GPT2 chatbot, showcased promising results in predicting mental health statuses. Logistic regression outperformed other algorithms, achieving an 80% accuracy rate. The study's holistic approach offers a significant stride in the realm of mental health research.

Study 6 A novel technique, "Sample Gradient Descent," was introduced, enhancing computational efficiency in comparison to traditional batch gradient descent. Leveraging PCA for representative data selection and optimizing hyperparameters led to faster convergence rates. The approach showcased adaptability and effectiveness across diverse tasks and datasets, highlighting the robustness and versatility of large-scale language models.

Overall, the results of these studies provide valuable insights into various aspects of NLP and demonstrate the potential of machine learning techniques in solving real-world problems.

13. Future Work

- Further investigation of the relationship between language and sentiment: This thesis explored the impact of language on sentiment analysis, but future research could delve deeper into the relationship between language and sentiment. This could include exploring the effect of linguistic nuances, idioms, and cultural differences on sentiment analysis.

- Developing more robust machine learning models: Although the study identified logistic regression as the most proficient machine learning algorithm for forecasting mental health status, there exists an opportunity for future investigations to delve into the possibilities presented by alternative machine learning techniques like deep learning, ensemble methods, and neural networks. Such exploration holds the potential to yield sentiment analysis models that are not only more precise but also more dependable.

- Exploring the ethical and social implications of using social media data for sentiment analysis: As the use of social media data for sentiment analysis becomes more widespread, there is a need to examine the ethical and social implications of this practice. Future research could investigate the potential privacy concerns, bias, and discrimination that could arise from the use of social media data for sentiment analysis.

- Investigating the potential of sentiment analysis for other applications: While this thesis focused on the use of sentiment analysis for mental health prediction, there are other potential applications of sentiment analysis that could be explored. For example, sentiment analysis could be used to improve customer service, political analysis, or market research.

- Developing new techniques for sentiment analysis: This thesis introduced a novel approach to gradient descent termed "Sample Gradient Descent" for sentiment analysis. Subsequent research endeavors could investigate the efficacy of alternative optimization techniques, such as stochastic gradient descent or Bayesian optimization, in the realm of sentiment analysis. Such exploration could contribute valuable insights to the field.

- Developing more effective interventions for mental health: This thesis identified a potential relationship between hate speech and mental health, particularly in individuals with anxiety and bipolar disorders. The creation of interventions to lessen the detrimental effects of hate speech on mental health could be the subject of future study.

- Expanding the scope of sentiment analysis to other languages and cultures: This thesis focused on sentiment analysis for English language social media text. Future research could explore the potential of sentiment analysis for other languages and cultures, which could require the development of new techniques and models.

- Advancing Large Language Models (LLMs) for Multimodal Understanding: While this study has primarily focused on textual data, future research could explore the integration of Large Language Models (LLMs) with multimodal data, including images, videos, and audio. Developing LLMs capable of understanding and generating content across multiple modalities could revolutionize fields such as multimedia content creation, accessibility technologies, and human-computer interaction. Investigating techniques for seamless fusion of textual and non-textual information within LLMs could pave the way for more comprehensive and context-aware artificial intelligence systems.

Author Publications

[1] Ganie, A. G., & Dadvandipour, S. (2021). Sentiment analysis on the effect of trending source less News: special reference to the recent death of an Indian actor. In *Artificial Intelligence and Sustainable Computing for Smart City: First International Conference, AIS2C2 2021, Greater Noida, India, March 22–23, 2021, Revised Selected Papers 1* (pp. 3-16). Springer International Publishing.
[2] Ganie, A. G., & Dadvandipour, S. (2022). Identification of online harassment using ensemble fine-tuned pre-trained Bert. *Pollack Periodica*.
[3] Aadil, G., & Samad, D. (2022). Traditional or deep learning for sentiment analysis: A review. *Multidiszciplináris Tudományok*, *12*(1), 3-12.
[4] MACHINE, M. U. (2021). Detection of semantic obsessive text in multimedia using machine and deep learning techniques and algorithms. *Journal of Theoretical and Applied Information Technology*, *99*(11).
[5] Ganie, A. G. (2023). Presence of informal language, such as emoticons, hashtags, and slang, impact the performance of sentiment analysis models on social media text?. *arXiv preprint arXiv:2301.12303*.
[6] Ganie, A. G. (2021). Private network optimization. *Multidiszciplináris Tudományok*, *11*(4), 248-254.
[7] Ganie, A., & Dadvandipour, S. (2023). Social Media Posts as a Window into Mental Health: A Machine Learning Approach.
[8] Dadvandipour, S., & Ganie, A. G. (2022). An Approach to Implementation of Autoencoders in Intelligent Vehicles. In *Vehicle and Automotive Engineering* (pp. 3-10). Cham: Springer International Publishing.
[9] Aadil, G., & Samad, D. (2022). Traditional or deep learning for sentiment analysis: A review. *Multidiszciplináris Tudományok*, *12*(1), 3-12.
[10] Aadil, G. G., & Samad, D. (2023). Exploring the Impact of Informal Language on Sentiment Analysis Models for Social Media Text Using Convolutional Neural Networks. *Multidiszciplináris Tudományok*, *13*(1), 244-254.
[11] Ganie, A. G., & Dadvandipour, S. (2023). From big data to smart data: a sample gradient descent approach for machine learning. *Journal of Big Data*, *10*(1), 162.
[12] Ganie, A. G., & Dadvandipour, S. (2023). From big data to smart data: a sample gradient descent approach for machine learning. Journal of Big Data, 10(1), 162.
[13] Lone, M. A., Khanday, O. M., & Gani, A. G. (2021). A Survey on Robot Behavior and Distance Estimation in IndoorGML Maps Implementation. *American Journal of Electronics & Communication*, *1*(3), 1-7.

References

1. Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*. https://doi.org/10.1561/1500000011
2. Liu, B. (2012). Sentiment Analysis and Opinion Mining Mining. Synthesis Lectures on Human Language Technologies. [Draft]. In *Synthesis Lectures on Human Language Technologies*.
3. Hu, M., & Liu, B. (2004). Mining and summarizing customer reviews. *KDD-2004 - Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. https://doi.org/10.1145/1014052.1014073
4. Camras, L., & Plutchik, R. (1980). Emotion: A Psychoevolutionary Synthesis. *The American Journal of Psychology*. https://doi.org/10.2307/1422394
5. Baccianella, S., Esuli, A., & Sebastiani, F. (2008). S ENTI W ORD N ET 3 . 0 : An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining. *Proceedings of the 7th Conference on Language Resources and Evaluation LREC10*.
6. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*.
7. Ganie, A. G., & Dadvandipour, S. (2021). Sentiment Analysis on the Effect of Trending Source Less News: Special Reference to the Recent Death of an Indian Actor. *Communications in Computer and Information Science*. https://doi.org/10.1007/978-3-030-82322-1_1
8. Hill, J. (2016). The Impact of Emojis and Emoticons on Online Consumer Reviews , Perceived Company Response Quality , Brand Relationship , and Purchase Intent . *University of South Florida Scholar Commons Graduate*.
9. Herring, S. C., Stein, D., & Virtanen, T. (2013). Introduction to the pragmatics of computer-mediated communication. In *Pragmatics of Computer-Mediated Communication*. https://doi.org/10.1515/9783110214468.3
10. Isnain, A. R., Sihabuddin, A., & Suyanto, Y. (2020). Bidirectional Long Short Term Memory Method and Word2vec Extraction Approach for Hate Speech Detection. *IJCCS (Indonesian Journal of Computing and Cybernetics Systems)*. https://doi.org/10.22146/ijccs.51743
11. Kim, J., Lee, J., Park, E., & Han, J. (2020). A deep learning model for detecting mental illness from user content on social media. *Scientific Reports*, *10*(1), 1–6.
12. Krishnamoorthy, S., Dua, A., & Gupta, S. (2021). Role of emerging technologies in future IoT-driven Healthcare 4.0 technologies: a survey, current challenges and future directions. *Journal of Ambient Intelligence and Humanized Computing*. https://doi.org/10.1007/s12652-021-03302-w
13. Liu, C., Fang, F., Lin, X., Cai, T., Tan, X., Liu, J., & Lu, X. (2021). Improving sentiment analysis accuracy with emoji embedding. *Journal of Safety Science and Resilience*, *2*(4), 246–252.
14. Modha, S., Majumder, P., & Mandl, T. (2018). Filtering Aggression from the Multilingual Social Media Feed. *COLING 2018 - 1st Workshop on Trolling, Aggression and Cyberbullying, TRAC 2018 - Proceedings of the Workshop*.
15. Nobata, C., Tetreault, J., Thomas, A., Mehdad, Y., & Chang, Y. (2016). Abusive language detection in online user content. *25th International World Wide Web Conference, WWW 2016*. https://doi.org/10.1145/2872427.2883062

16. Poria, S., Cambria, E., & Gelbukh, A. (2016). Aspect extraction for opinion mining with a deep convolutional neural network. *Knowledge-Based Systems*. https://doi.org/10.1016/j.knosys.2016.06.009
17. Proceedings - 2018 IEEE International Conference on Big Data, Big Data 2018. (2019). In *Proceedings - 2018 IEEE International Conference on Big Data, Big Data 2018*.
18. Reece, A. G., & Danforth, C. M. (2017). Instagram photos reveal predictive markers of depression. *EPJ Data Science*. https://doi.org/10.1140/epjds/s13688-017-0110-z
19. Severyn, A., & Moschitti, A. (2015a). Learning to Rank Short Text Pairs with Convolutional Deep Neural Networks Categories and Subject Descriptors. *Sigir*.
20. Shing, H. C., Nair, S., Zirikly, A., Friedenberg, M., Daumé, H., & Resnik, P. (2018). Expert, crowdsourced, and machine assessment of suicide risk via online postings. *Proceedings of the 5th Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic, CLPsych 2018 at the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies,*. https://doi.org/10.18653/v1/w18-0603
21. Singh, V., Varshney, A., Akhtar, S. S., Vijay, D., & Shrivastava, M. (2018). Aggression Detection on Social Media Text Using Deep Neural Networks. *2nd Workshop on Abusive Language Online - Proceedings of the Workshop, Co-Located with EMNLP 2018*. https://doi.org/10.18653/v1/w18-5106
22. Sutskever, I., Martens, J., Dahl, G., & Hinton, G. (2013). On the importance of initialization and momentum in deep learning. *30th International Conference on Machine Learning, ICML 2013*.
23. Tsugawa, S., Kikuchi, Y., Kishino, F., Nakajima, K., Itoh, Y., & Ohsaki, H. (2015). Recognizing depression from twitter activity. *Conference on Human Factors in Computing Systems - Proceedings*. https://doi.org/10.1145/2702123.2702280