Enhancing Emotion Classification with a Hybrid BERT and CNN Architecture

Orestis Papadimitriou*, Athanasios Kanavos*, Gerasimos Vonitsanos†,
Manolis Maragoudakis‡ and Phivos Mylonas§
*Department of Information and Communication Systems Engineering
University of the Aegean, Samos, Greece
{icsdd20016, icsdd20017}@icsd.aegean.gr
†Computer Engineering and Informatics Department
University of Patras, Patras, Greece
mvonitsanos@ceid.upatras.gr
‡Department of Informatics
Ionian University, Corfu, Greece
mmarag@ionio.gr
§Department of Informatics and Computer Engineering
University of West Attica, Athens, Greece
mylonasf@uniwa.gr

***Abstract*—In computational linguistics, effectively encoding and systematizing emotional expressions in language is a significant challenge. Existing machine learning (ML) models for text analysis often only recognize primary emotional states such as anger, happiness, and sadness, missing the more nuanced spectrum of human emotions. These models frequently overlook the complex semantic interplays and fail to capture the full diversity of emotional expressions, focusing instead on simplistic categorization. This paper proposes a robust ML framework that enhances emotion classification by discriminating among nuanced emotion categories—Anger, Joy, and Fear. Leveraging a sophisticated combination of Convolutional Neural Networks (CNNs) and the Bidirectional Encoder Representations from Transformers (BERT) architecture, our framework demonstrates exceptional accuracy in emotion detection. Evaluating a diverse dataset of text samples, each meticulously tagged with its expressed emotion, confirms the model's superior performance in recognizing and classifying a broad range of emotional states.**

***Index Terms*—Emotion Classification, Computational Linguistics, Machine Learning, Social Media Text Analysis, Data Analysis**

I. Introduction

Natural Language Processing (NLP) is pivotal in decoding human emotions from text, a task underscored by the daily influx of around 500 million tweets on platforms like Twitter [46]. The overwhelming volume of data challenges human capacity for analysis, necessitating sophisticated automated systems to interpret complex emotions beyond primary states like happiness or sadness. Despite advancements, the nuanced nature of human language, including sarcasm, irony, and cultural expression variations, poses significant challenges [31], [48].

979-8-3503-2771-7/23/$31.00 © 2023 IEEE

Traditional sentiment analysis tools often rely on categorizing text into primary emotional states and need help with social media language's complexity and informal nature [19], [20]. Furthermore, these tools cannot typically adapt to multilingual contexts, each with unique linguistic structures [21], [36]. The dynamic and evolving nature of social media demands more sophisticated techniques to grasp the explicit and implicit cues like tone, context, and cultural nuances.

Recent advances in machine learning, mainly through Convolutional Neural Networks (CNNs), have significantly improved emotion detection. However, challenges such as extensive training times and diminished accuracy in complex scenarios remain. Machine learning, especially supervised learning, plays a crucial role by training algorithms on labeled datasets to develop models that proficiently categorize new instances [5].

Emotions are critical in understanding interpersonal communications. Beyond the simplistic positive, negative, or neutral categorizations, it is essential to discern specific emotions like joy, sadness, or anger, which enrich our understanding of human interactions [30]. Most existing methods classify emotions into broad categories proposed by theorists like Plutchik, which do not encompass the entire emotional spectrum, such as guilt, shame, or pride [37], [38], [42].

Integrating linguistics, psychology, and social sciences knowledge is crucial in crafting sentiment analysis tools that transcend mere technological solutions [45]. These tools must understand and adapt to linguistic and cultural variations across different societies [26]. The evolution of sentiment analysis is likely to emphasize adaptable systems that interpret and learn from the changing patterns of online communication.

Groundbreaking developments in deep learning, particularly with the advent of models like BERT (Bidirectional Encoder

Representations from Transformers), have demonstrated the potential to transcend the limitations of traditional sentiment analysis methods [13]. BERT's use of attention mechanisms allows for a nuanced understanding of contextual relationships between words, thus improving the detection of complex emotional expressions [12]. However, the practical application of such models faces challenges related to computational demands and model interpretability [28].

Hybrid models that combine the pattern-recognition capabilities of CNNs with the context-sensitive attributes of Transformer-based architectures like BERT offer promising solutions [2], [16], [43], [47]. These models benefit from both local pattern recognition and the understanding of broader semantic contexts, potentially enhancing the accuracy and efficiency of emotion detection systems [8].

The remainder of this paper is organized as follows: Section II delves into the related work, exploring existing methodologies and advancements that have shaped current approaches to emotion classification in text analysis. Section III details the methodology foundations, describing the specific architectures and configurations of BERT and CNN used in this study. Section IV evaluates the effectiveness of these methods through a series of experiments, assessing the impact of various parameters on model performance. Finally, Section V concludes the paper with a discussion of the findings, implications for the field, and potential directions for future research. This structured approach ensures a thorough understanding of the challenges and innovations in emotion detection from text, aiming to contribute significantly to the advancements in natural language processing and machine learning.

II. Related Work

The field of sentiment analysis has significantly advanced our understanding of human emotions conveyed through text, particularly in the face of the exponential growth of data on social platforms. This section reviews the evolution of neural network applications in sentiment analysis, emphasizing various methodologies and their efficacy in capturing the subtleties of expressed emotions [7], [44].

Recent advancements in deep learning have considerably enhanced the capabilities of traditional machine learning techniques in sentiment analysis [15], [32], [40]. Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) have become pivotal for their high accuracy in recognition tasks [24], [33]. CNNs, notable for their layered architecture, excel at processing spatial relationships in data, using functions such as ReLU to introduce non-linearity, which is crucial for analyzing complex, real-world information [23], [22], [34], [35]. Their versatility in managing diverse data types makes them exceptionally effective across various social media platforms.

However, sentiment analysis tools frequently encounter issues with consistency across diverse datasets and languages, highlighting the need for more robust models [39]. To address these challenges, keyword-based strategies have evolved, with new databases enhancing the detection of emotion-related phrases [41]. Additionally, learning-based methods employing a variety of classifiers have shown promise, particularly when contrasting the efficacy of traditional models against more sophisticated machine learning techniques such as KNN, XGBoost, and SVM [6].

Integrating symbolic knowledge into machine learning models has significantly improved their ability to recognize complex emotions in conversations, markedly outperforming baseline models [27]. Adopting pre-trained language models (PTLMs) underscores the importance of transfer learning in enhancing model adaptability across different emotional contexts and languages [3]. This adaptability is crucial for applications in diverse, multilingual environments where emotional expressions vary widely.

Moreover, neurosymbolic AI systems that merge subsymbolic and symbolic approaches have been developed to interpret better and analyze sentiments [10]. Combining traditional and modern techniques, these systems offer a holistic approach to sentiment analysis. Cross-lingual techniques further utilize emotion annotations from one language to classify emotions in texts of another without direct translation, leveraging frameworks like Ekman's emotional categories to bridge cultural differences [9].

What is more, Bi-directional LSTM (BiLSTM) networks represent a significant advancement, offering bidirectional text processing and the capability to manage texts of varying lengths through advanced learning strategies like sent2affect. Initially trained for sentiment analysis, this method is then adapted for nuanced emotion detection, achieving performance comparable to established machine learning algorithms [25]. The flexibility of BiLSTMs in processing the temporal dynamics of language makes them especially suitable for analyzing conversational data where context and sequence of expressions are crucial.

The exploration of sentiment analysis has benefited from a cross-pollination of techniques, from profound learning innovations to hybrid models that blend different computational architectures. As this field progresses, the focus shifts towards models that can seamlessly integrate into diverse linguistic and cultural frameworks, enhancing their utility and accuracy in real-world applications. Finally, the continued refinement of these technologies is imperative for developing practical sentiment analysis tools across various platforms and languages capable of adapting to the evolving nuances of human communication.

III. Methodology Foundations

This section outlines the foundational technologies and architectures employed in our study to address the challenge of emotion classification in natural language processing. We utilize advanced models such as BERT (Bidirectional Encoder Representations from Transformers) and Convolutional Neural Networks (CNNs), combining their respective strengths to enhance the understanding and processing of language. BERT provides deep contextual insights into language structure, which is complemented by CNNs' spatial feature extraction capabilities, creating a robust framework for accurate and efficient text analysis.

A. BERT Architecture

In this research, we implemented a BERT architecture to facilitate emotion classification in natural language processing. Since its introduction, BERT has revolutionized NLP by enabling models to understand the contextual meanings of words within text sequences [12].

BERT is uniquely designed for pre-training deep bidirectional representations by conditioning on both left and right contexts in all layers. Consequently, the pre-trained BERT model can be fine-tuned with just one additional output layer to accommodate a wide range of tasks, such as question answering and language inference, without extensive modifications to its architecture.

A cornerstone of BERT's innovation is its self-attention mechanism, described as follows:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{1}$$

where Q, K, and V represent the query, key, and value vectors, respectively, and d_k is the dimensionality of the vital vector, which scales the dot products.

Each layer of BERT transforms its input using the following operations:

$$\text{LayerOutput} = \text{LayerNorm}(x + \text{Sublayer}(x)) \tag{2}$$

where $\text{Sublayer}(x)$ is the function implemented by the sublayers of the model (either multi-head self-attention or fully connected feed-forward neural networks), and LayerNorm is the layer normalization function.

BERT also employs Masked Language Modeling (MLM) during pre-training, where it randomly masks some of the tokens from the input, and the objective is to predict the original vocabulary id of the masked word based on its context. The loss function for MLM is given by:

$$\text{Loss}_{\text{MLM}} = -\log(\text{softmax}(x_{\text{masked}})_{\text{original\_id}}) \tag{3}$$

In our implementation, the model follows a structured pipeline for data pre-processing, training, validation, and testing [14]. The dataset is divided into training, validation, and test sets. The model is initialized with an appropriate optimizer and deployed on the optimal hardware environment (GPU or CPU). DataLoader objects manage data batching and shuffling, facilitating efficient input handling during training.

B. CNN Enhancement for Text Analysis

The BERT model employed in our study is the foundation for a robust CNN architecture explicitly designed for text analysis. BERT transforms raw text into a format comprehensible to the model, breaking it into smaller tokens. These tokens are then processed through multiple model layers, enabling BERT to learn and interpret the relationships between words within a sentence. Given its pre-trained nature, BERT understands language dynamics but requires fine-tuning to adapt to specific tasks like emotion classification or sentiment analysis.

- **Pre-trained BERT model:**
  - Utilizes the 'bert-base-uncased' version, which is trained on a lower-cased English text corpus and is adept at general language understanding.
  - Customized for text classification, the model adapts to the nuances of the specific task by adjusting its final layers to correspond to the number of categories in the label mapping.
- **AdamW optimizer:**
  - Employs AdamW, an advanced optimizer that modifies the model parameters during training more efficiently than traditional stochastic gradient descent.
  - Adjusts the learning rate ('LR') and epsilon ('EPSILON') settings to fine-tune the convergence behavior of the model, ensuring optimal updates with minimal risk of diverging.
- **Device selection:**
  - Automatically detects hardware capabilities, prioritizing GPU usage for its superior processing speed, but can revert to CPU if a GPU is unavailable.
- **Model deployment to the device:**
  - Deploys the BERT model to the selected hardware (GPU or CPU) to leverage computational efficiency during training and inference phases.

During training, the model undergoes rigorous testing against a dataset divided into training, validation, and test subsets. This segmentation ensures that the model learns effectively and validates its learning against unseen data, which is crucial for assessing its generalization capability. The AdamW optimizer facilitates subtle adjustments in the model's weights and biases, directly influencing its ability to make accurate predictions.

Adapting the BERT model within a CNN framework involves integrating convolutional layers that process the tokenized input. These layers capture local dependencies and critical features in text data, enhancing the model's ability to discern subtle emotional cues and semantic patterns. Convolutional layers apply filters to the input, extracting essential features for the classification task. Precisely, each convolutional layer is followed by a non-linear activation function, typically ReLU, introducing non-linearity that enables the model to learn complex patterns.

The strategic deployment of the model to an appropriate computational device is essential for managing resource utilization and achieving computational efficiency. The choice between GPU and CPU can significantly affect training times and overall model performance, making device selection a critical aspect of model configuration.

By leveraging the sophisticated capabilities of both BERT's pre-trained knowledge and CNN's architectural advantages, this hybrid approach significantly enhances the accuracy and efficiency of text classification tasks, promising substantial advancements in emotion detection and analysis.

IV. Evaluation

This section provides a detailed evaluation of the proposed model using the Emotion Classification dataset. We assess the model's performance through various experiments focusing on the impact of batch size adjustments on training dynamics, accuracy, and loss. This comprehensive evaluation helps validate the model's effectiveness in classifying emotions from text and offers insights into the scalability and efficiency of our approach under different computational loads.

A. Dataset

The Emotion Classification dataset [1] serves as the cornerstone for our evaluation of the proposed BERT and CNN-enhanced text analysis architecture. This dataset is explicitly curated to advance research in natural language processing and emotion analysis, containing various text samples, each annotated with an expressed emotion. These annotations cover a comprehensive range of emotions, from joy and enthusiasm to anger and sadness, making it an ideal resource for training and validating our emotion classification model.

Each text sample in the dataset is tagged to reflect the expressed emotion, providing a rich basis for analyzing the effectiveness of different NLP models in recognizing and categorizing emotional content. The dataset's diversity in emotional range and linguistic variations presents a challenging yet realistic environment for testing the robustness of our methodologies.

B. Model Performance Evaluation

The effects of batch size on the performance of our NLP model are explored through systematic variation during training. These experiments assess impacts on training dynamics, including training and validation accuracies, losses, and epoch training times, detailed in Table I and visually represented in Figure 1.

With a batch size of 16, the model achieves relatively high training and validation accuracies, indicating effective learning. Low training and validation losses support these results, suggesting that the model is generalizing well to unseen data. However, the training time per epoch is relatively high, averaging 80.7 seconds, which may affect scalability and training efficiency.

Increasing the batch size to 32 still yields good performance in training and validation accuracies. Although the losses are slightly higher than with a batch size of 16, they remain low, indicating a slight decrease in model generalization. The training time per epoch decreases to an average of 71 seconds, suggesting improved training efficiency without significantly compromising performance.

With a batch size of 64, there is a noticeable decrease in training and validation accuracies, as the model struggles more with generalization. This is reflected in higher losses compared to smaller batch sizes. The training time per epoch reduces further to an average of 69 seconds, which may be advantageous for scenarios with larger datasets or when computational resources are a limiting factor.

At a batch size of 128, the model exhibits the lowest training and validation accuracies among the tested configurations. The losses are higher, indicating poorer model generalization across the board. While the training time per epoch is the shortest, averaging 66.3 seconds, this comes at the cost of decreased model performance.

In summary, smaller batch sizes (16 and 32) generally lead to better model performance and generalization at the expense of longer training times per epoch. Conversely, larger batch sizes (64 and 128) yield faster training times but may suffer from decreased model performance and generalization. The choice of batch size should consider the trade-offs between training efficiency and model effectiveness based on the specific requirements and constraints of the task at hand.

V. Conclusions and Future Work

This study has successfully demonstrated the efficacy of integrating BERT with a convolutional neural network (CNN) architecture for emotion classification in natural language processing. Our findings indicate that the hybrid model not only leverages the deep contextual understanding of BERT but also harnesses the spatial feature extraction capabilities of CNNs, resulting in a robust framework for text analysis. Smaller batch sizes consistently yielded higher accuracies and lower losses throughout the experiments, suggesting that the model could effectively learn and generalize from the training data.

However, smaller batch sizes showed improved performance metrics, resulting in longer training times. This trade-off highlights an essential consideration for deploying such models in real-world applications where computational resources and time constraints are factors. The results also point to the significant impact that batch size has on training dynamics, emphasizing the need for careful parameter tuning to balance efficiency and effectiveness.

Several avenues appear promising for future work. First, further exploration into more dynamic batch sizing techniques could be conducted to optimize training efficiency without compromising model performance. Adaptive batching strategies, which adjust the batch size based on training progress, might improve training time and model accuracy [4]. Additionally, expanding the model's capability to include more nuanced emotional states or incorporating multilingual datasets could enhance its applicability and robustness across different linguistic and cultural contexts [11], [17], [18].

Moreover, exploring alternative architectures or newer NLP models, such as Transformer-based models like GPT-3 or RoBERTa, might offer further improvements in both performance and efficiency. These models provide advancements in handling longer dependencies and could be tested for effectiveness in emotion classification tasks. Implementing such models could also be paralleled with efforts to reduce computational demands, perhaps through more efficient model pruning techniques or leveraging newer, more efficient hardware technologies.

TABLE I
EXPERIMENTAL EVALUATION OF MODEL PERFORMANCE ACROSS DIFFERENT BATCH SIZES

Epochs	Loss	Accuracy	Time	Loss	Accuracy	Time
	Batch Size = 16			**Batch Size = 32**		
1	0.6287	0.8979	77	0.7870	0.8842	69
5	0.0260	0.9432	83	0.0327	0.9484	72
10	0.0143	0.9547	82	0.0122	0.9495	72
	Batch Size = 64			**Batch Size = 128**		
1	0.9870	0.7000	69	1.056	0.5579	66
5	0.0608	0.9400	69	0.1550	0.9232	67
10	0.0186	0.9432	69	0.0325	0.9421	66

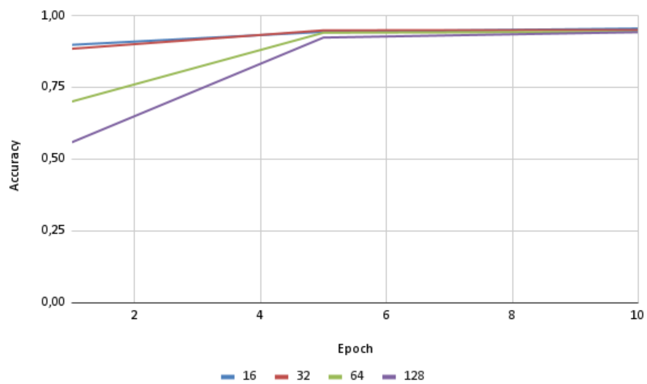

(a) Accuracy

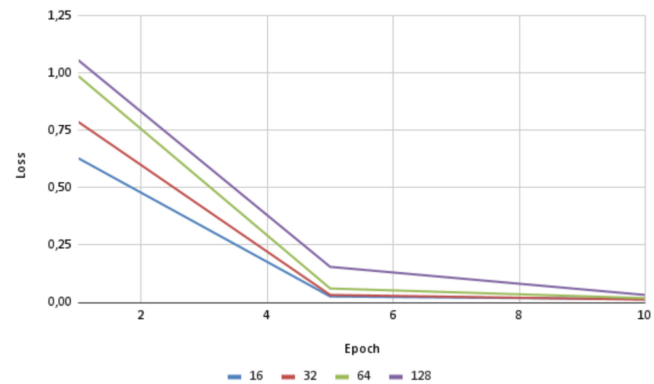

(b) Loss

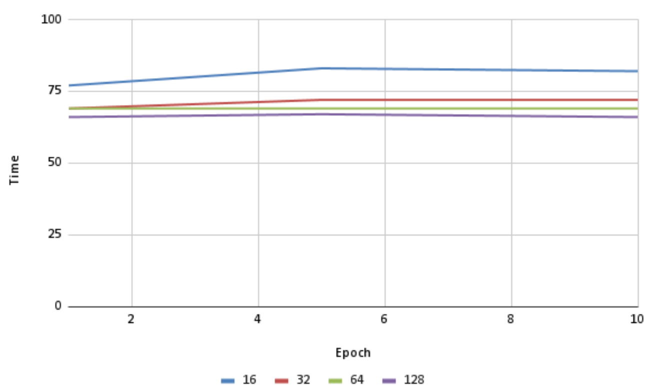

(c) Time

Fig. 1. Accuracy, Loss, and Time Analysis for Different Batch Sizes

Ultimately, this research opens up several pathways for enhancing text analysis technologies in NLP, paving the way for more sensitive and accurate emotion detection tools, which are crucial for various applications, from customer service to therapeutic aids [29].

REFERENCES

[1] Emotion dataset. https://www.kaggle.com/datasets/abdallahwagih/emotion-dataset, online; accessed on 11 August 2024

[2] Abas, A.R., Elhenawy, I., Zidan, M., Othman, M.: Bert-cnn: A deep learning model for detecting emotions from text. Computers, Materials & Continua **71**(2) (2022)

[3] Ahmad, Z., Jindal, R., Ekbal, A., Bhattachharyya, P.: Borrow from rich cousin: Transfer learning for emotion detection using cross lingual embedding. Expert Systems with Applications **139** (2020)

[4] Alexopoulos, A., Drakopoulos, G., Kanavos, A., Mylonas, P., Vonitsanos, G.: Two-step classification with SVD preprocessing of distributed massive datasets in apache spark. Algorithms **13**(3), 71 (2020)

[5] Alexopoulos, A., Drakopoulos, G., Kanavos, A., Sioutas, S., Vonitsanos, G.: Parametric evaluation of collaborative filtering over apache spark. In: 5th South-East Europe Design Automation, Computer Engineering, Computer Networks and Social Media Conference (SEEDA-CECNSM). pp. 1–8. IEEE (2020)

[6] Alotaibi, F.M.: Classifying text-based emotions using logistic regression. VAWKUM Transactions on Computer Sciences **7**(1), 31–37 (2019)

[7] Andrea, A.D., Ferri, F., Grifoni, P., Guzzo, T.: Approaches, tools and applications for sentiment analysis implementation. International Journal of Computer Applications **125**(3) (2015)

[8] Arriaga, O., Valdenegro-Toro, M., Plöger, P.: Real-time convolutional neural networks for emotion and gender classification. CoRR **abs/1710.07557** (2017)
[9] Becker, K., Moreira, V.P., dos Santos, A.G.L.: Multilingual emotion classification using supervised learning: Comparative experiments. Information Processing and Management **53**(3), 684–704 (2017)
[10] Cambria, E., Liu, Q., Decherchi, S., Xing, F., Kwok, K.: Senticnet 7: A commonsense-based neurosymbolic AI framework for explainable sentiment analysis. In: 13th Language Resources and Evaluation Conference (LREC). pp. 3829–3839. European Language Resources Association (2022)
[11] Das, A., Gambäck, B.: Sentimantics: Conceptual spaces for lexical sentiment polarity representation with contextuality. In: 3rd Workshop in Computational Approaches to Subjectivity and Sentiment Analysis (WASSA@ACL). pp. 38–46. The Association for Computer Linguistics (2012)
[12] Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT). pp. 4171–4186. Association for Computational Linguistics (2019)
[13] Do, H.H., Prasad, P.W.C., Maag, A., Alsadoon, A.: Deep learning for aspect-based sentiment analysis: A comparative review. Expert Systems with Applications **118**, 272–299 (2019)
[14] Dritsas, E., Vonitsanos, G., Livieris, I.E., Kanavos, A., Ilias, A., Makris, C., Tsakalidis, A.K.: Pre-processing framework for twitter sentiment classification. In: Artificial Intelligence Applications and Innovations (AIAI). vol. 560, pp. 138–149. Springer (2019)
[15] Giatsoglou, M., Vozalis, M.G., Diamantaras, K.I., Vakali, A., Sarigiannidis, G., Chatzisavvas, K.C.: Sentiment analysis leveraging emotions and word embeddings. Expert Systems with Applications **69**, 214–224 (2017)
[16] Jaiswal, S., Nandi, G.C.: Robust real-time emotion detection system using CNN architecture. Neural Computing and Applications **32**(15), 11253–11262 (2020)
[17] Kafeza, E., Kanavos, A., Makris, C., Pispirigos, G., Vikatos, P.: T-PCCE: twitter personality based communicative communities extraction system for big data. IEEE Transactions on Knowledge and Data Engineering **32**(8), 1625–1638 (2020)
[18] Kafeza, E., Kanavos, A., Makris, C., Vikatos, P.: T-PICE: twitter personality based influential communities extraction system. In: International Congress on Big Data. pp. 212–219. IEEE Computer Society (2014)
[19] Kanavos, A., Perikos, I., Hatzilygeroudis, I., Tsakalidis, A.K.: Emotional community detection in social networks. Computers & Electrical Engineering **65**, 449–460 (2018)
[20] Kanavos, A., Vonitsanos, G., Mohasseb, A., Mylonas, P.: An entropy-based evaluation for sentiment analysis of stock market prices using twitter data. In: 15th International Workshop on Semantic and Social Media Adaptation and Personalization (SMAP). pp. 1–7. IEEE (2020)
[21] Kanavos, A., Kolovos, E., Papadimitriou, O., Maragoudakis, M.: Breast cancer classification of histopathological images using deep convolutional neural networks. In: 7th South-East Europe Design Automation, Computer Engineering, Computer Networks and Social Media Conference (SEEDA-CECNSM). pp. 1–6. IEEE (2022)
[22] Kanavos, A., Papadimitriou, O., Kaponis, A., Maragoudakis, M.: Enhancing disease diagnosis: A cnn-based approach for automated white blood cell classification. In: IEEE International Conference on Big Data. pp. 4606–4613 (2023)
[23] Kanavos, A., Papadimitriou, O., Maragoudakis, M.: Enhancing COVID-19 diagnosis from chest x-ray images using deep convolutional neural networks. In: 18th International Workshop on Semantic and Social Media Adaptation and Personalization (SMAP). pp. 1–6. IEEE (2023)
[24] Kanavos, A., Papadimitriou, O., Mylonas, P., Maragoudakis, M.: Enhancing sign language recognition using deep convolutional neural networks. In: 14th International Conference on Information, Intelligence, Systems & Applications (IISA). pp. 1–4. IEEE (2023)
[25] Kratzwald, B., Ilic, S., Kraus, M., Feuerriegel, S., Prendinger, H.: Deep learning for affective computing: Text-based emotion recognition in decision support. Decision Support Systems **115**, 24–35 (2018)
[26] Leary, M.R.: Motivational and emotional aspects of the self. Annual Review of Psychology **58**, 317–344 (2007)
[27] Li, W., Zhu, L., Mao, R., Cambria, E.: SKIER: A symbolic knowledge integrated model for conversational emotion recognition. In: 37th AAAI Conference on Artificial Intelligence (AAAI). pp. 13121–13129. AAAI Press (2023)
[28] Makiuchi, M.R., Warnita, T., Uto, K., Shinoda, K.: Multimodal fusion of BERT-CNN and gated CNN representations for depression detection. In: 9th International on Audio/Visual Emotion Challenge and Workshop (AVEC@MM). pp. 55–63. ACM (2019)
[29] Mohasseb, A., Kanavos, A.: Grammar-based question classification using ensemble learning algorithms. In: 18th International Conference on Web Information Systems and Technologies (WEBIST). Lecture Notes in Business Information Processing, vol. 494, pp. 84–97. Springer (2022)
[30] Munezero, M., Montero, C.S., Sutinen, E., Pajunen, J.: Are they different? affect, feeling, emotion, sentiment, and opinion detection in text. IEEE Transactions on Affective Computing **5**(2), 101–111 (2014)
[31] Oh, S.J., Lee, J., Kim, D.K.: The design of CNN architectures for optimal six basic emotion classification using multiple physiological signals. Sensors **20**(3), 866 (2020)
[32] Pang, B., Lee, L., Vaithyanathan, S.: Thumbs up? sentiment classification using machine learning techniques. In: Conference on Empirical Methods in Natural Language Processing (EMNLP). pp. 79–86 (2002)
[33] Papadimitriou, O., Kanavos, A., Maragoudakis, M.: Automated pneumonia detection from chest x-ray images using deep convolutional neural networks. In: 14th International Conference on Information, Intelligence, Systems & Applications (IISA). pp. 1–4. IEEE (2023)
[34] Papadimitriou, O., Kanavos, A., Maragoudakis, M., Gerogiannis, V.C.: Chess piece recognition using deep convolutional neural networks. In: 4th Symposium on Pattern Recognition and Applications (SPRA). vol. 13162, p. 1316202 (2024)
[35] Papadimitriou, O., Kanavos, A., Mylonas, P., Maragoudakis, M.: Advancing weather image classification using deep convolutional neural networks. In: 18th International Workshop on Semantic and Social Media Adaptation and Personalization (SMAP). pp. 1–6. IEEE (2023)
[36] Papadimitriou, O., Kanavos, A., Mylonas, P., Maragoudakis, M.: Classification of alzheimer's disease subjects from MRI using deep convolutional neural networks. In: 3rd International Conference on Novel & Intelligent Digital Systems (NiDS). Lecture Notes in Networks and Systems, vol. 784, pp. 277–286. Springer (2023)
[37] Plutchik, R.: A general psychoevolutionary theory of emotion. In: Theories of Emotion. pp. 3–33 (1980)
[38] Poria, S., Hazarika, D., Majumder, N., Naik, G., Cambria, E., Mihalcea, R.: MELD: A multimodal multi-party dataset for emotion recognition in conversations. CoRR **abs/1810.02508** (2018)
[39] Ribeiro, F.N., Araújo, M., Gonçalves, P., Benevenuto, F., Gonçalves, M.A.: A benchmark comparison of state-of-the-practice sentiment analysis methods. CoRR **abs/1512.01818** (2015)
[40] Saad, F.: Baseline evaluation: An empirical study of the performance of machine learning algorithms in short snippet sentiment analysis. In: 14th International Conference on Knowledge Management and Data-driven Business (I-KNOW). pp. 6:1–6:8. ACM (2014)
[41] Seal, D., Roy, U.K., Basak, R.: Sentence-level emotion detection from text based on semantic rules. In: Information and Communication Technology for Sustainable Development (ICT4SD). pp. 423–430 (2020)
[42] Strapparava, C., Mihalcea, R.: Semeval-2007 task 14: Affective text. In: 4th International Workshop on Semantic Evaluations (SemEval@ACL). pp. 70–74. The Association for Computer Linguistics (2007)
[43] Talaat, A.S.: Sentiment analysis classification system using hybrid BERT models. Journal of Big Data **10**(1), 110 (2023)
[44] Tanna, D., Dudhane, M., Sardar, A., Deshpande, K., Deshmukh, N.: Sentiment analysis on social media for emotion classification. In: 4th International Conference on Intelligent Computing and Control Systems (ICICCS). pp. 911–915. IEEE (2020)
[45] Vernikou, S., Lyras, A., Kanavos, A.: Multiclass sentiment analysis on covid-19-related tweets using deep learning models. Neural Computing and Applications **34**(22), 19615–19627 (2022)
[46] Vonitsanos, G., Kanavos, A., Mylonas, P.: Decoding gender on social networks: An in-depth analysis of language in online discussions using natural language processing and machine learning. In: IEEE International Conference on Big Data. pp. 4618–4625 (2023)
[47] Zhang, L., Wang, S., Liu, B.: Deep learning for sentiment analysis: A survey. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery **8**(4) (2018)
[48] Zoupanos, S., Kolovos, S., Kanavos, A., Papadimitriou, O., Maragoudakis, M.: Efficient comparison of sentence embeddings. In: 12th Hellenic Conference on Artificial Intelligence (SETN). pp. 11:1–11:6. ACM (2022)