

Goodness: Intelligent System for Delivering Positive News with Sentiment Analysis and Summarization

S. Rubin Bose^{1,*}, R. S. Gaayathri², J. Angelin Jeba³, Judy Flavia⁴, R. Regin⁵, P. Paramasivan⁶

^{1,2,4,5}Department of Computer Science and Engineering, SRM Institute of Science and Technology, Ramapuram, Chennai, Tamil Nadu, India.

³Department of Electronics and Communication Engineering, S.A. Engineering College, Chennai, Tamil Nadu, India.

⁶Department of Research and Development, Dhaanish Ahmed College of Engineering, Chennai, Tamil Nadu, India.
rubinbos@srmist.edu.in¹, gr4829@srmist.edu.in², angelinjeba@saec.ac.in³, judyflab@srmist.edu.in⁴, reginr@srmist.edu.in⁵, paramasivanchem@gmail.com⁶

Abstract: In the modern digital landscape, individuals are inundated with vast amounts of news content, much of which carries a negative tone, contributing to heightened stress, anxiety, and emotional distress. To address this issue, the “Goodness” paper presents an advanced web-based system that collects, summarizes, and analyzes daily news articles, providing sentiment-annotated summaries through state-of-the-art machine learning and NLP techniques. The system follows a multi-stage pipeline. News articles are gathered from RSS feeds, processed using feed parser and trafilatura, and summarized via a BART-based model that condenses lengthy articles while preserving key details. Sentiment analysis is implemented through a dual-model approach. TextBlob filters neutral articles, and then a Logistic Regression model trained on the NLTK movie reviews dataset classifies the rest as positive or negative. A user-friendly Flask-based web interface with sentiment-based colour-coded highlights and ngrok for accessibility and scalability displays the summarised text. The “Goodness” technology improves digital news consumption by combining sentiment-aware summarisation with a real-time interactive interface to help consumers discern emotional tones and reduce the psychological impact of negative news. This study shows that AI-driven sentiment analysis can improve digital media consumption and public awareness.

Keywords: Natural Language Processing; Sentiment Analysis; Text Summarization; Flask Web Application; Negative Content; Pre-Trained Language Models; Decision-Making; NLP Community.

Received on: 01/06/2024, **Revised on:** 29/07/2024, **Accepted on:** 18/09/2024, **Published on:** 03/12/2024

Journal Homepage: <https://www.fmdbpublish.com/user/journals/details/FTSCL>

DOI: <https://doi.org/10.69888/FTSCL.2024.000281>

Cite as: S. R. Bose, R. S. Gaayathri, J. A. Jeba, J. Flavia, R. Regin, and P. Paramasivan, “Goodness: Intelligent System for Delivering Positive News with Sentiment Analysis and Summarization,” *FMDB Transactions on Sustainable Computer Letters*, vol. 2, no. 4, pp. 217–231, 2024.

Copyright © 2024 S. R. Bose *et al.*, licensed to Fernando Martins De Bulhão (FMDB) Publishing Company. This is an open access article distributed under [CC BY-NC-SA 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/), which permits copying, remixing, redistributing, transformation, and building upon the material in any medium so long as the original work is properly cited.

1. Introduction

The “Goodness” project aims to enhance how people consume daily news by leveraging advanced AI tools for sentiment analysis and news summarization. In an age where the volume of information can be overwhelming, this project addresses the need for concise, digestible, and sentiment-aware content presentation. By analyzing news articles and displaying summarized versions alongside sentiment classifications, “Goodness” makes it easier for readers to grasp the essence of news stories and

*Corresponding author.

their emotional undertones without sifting through lengthy content. The core objective of “Goodness” is to provide a seamless experience for users to stay informed, reducing information overload while ensuring they understand the emotional context behind the news. This is accomplished by integrating a summarization model, BART (Bidirectional and Auto-Regressive Transformers), to generate accurate and concise news summaries. The project further utilizes TextBlob, a natural language processing library, in combination with Logistic Regression, to perform sentiment analysis classifying articles as positive, negative, or neutral. Sentiment analysis plays a crucial role in understanding how news stories might impact readers, as it deciphers the emotional weight the content conveys [11].

Including a sentiment classifier allows “Goodness” to highlight these aspects, enabling users to filter news based on their mood or preference. The neutral classification provided by TextBlob is particularly beneficial for identifying unbiased or factual content [12]. The project uses a clean, user-friendly interface created with Flask, a lightweight web framework for Python. Users are presented with a list of summarized news items colour-coded based on positive, neutral, or negative sentiment. Each summary is accompanied by a link to the full article for those who wish to delve deeper [13]. By combining sophisticated machine learning techniques with intuitive design, “Goodness” offers a solution that saves users time and provides them with an emotionally aware reading experience.

This research aims to explore the methodologies employed in “Goodness” and evaluate their effectiveness in producing a reliable and impactful sentiment-aware news summarization system [14]. Negative news often captures more attention than positive or neutral stories due to its emotional impact. This tendency, known as “negativity bias,” can significantly affect readers’ moods, leading to increased anxiety, stress, or pessimism. Constant exposure to negative content can foster a distorted view of reality, making the world appear more threatening than it is [15]. This influence can reduce well-being, spark fear, and fuel a cycle where negative news is more frequently sought and consumed. The “Goodness” project aims to mitigate this impact by classifying and presenting content with clear sentiment markers, helping readers navigate emotionally charged stories.

2. Literature Review

Manikandan and Priscila [1] undertake sentiment analysis (S.A.) on COVID-19 tweets using natural language processing (NLP) techniques. They employ Flair PyTorch (FP) and TextBlob to analyze sentiment in tweets, incorporating text and emojis. The ‘en-sentiment’ module in FP is utilized for tokenization and text classification, discerning sentiments as positive or negative. The study contrasts FP’s performance with TextBlob, underscoring the value of emojis in enhancing sentiment analysis. Pourkeyvan et al. [2] stress the significance of early intervention in mental disorders. Their study uses social media data and pre-trained language models to predict symptoms, comparing four BERT models from Hugging Face with traditional methods. The results indicate superior performance of the new models, achieving up to 97% accuracy, highlighting the potential of social media for mental health screening even with limited data. The authors propose leveraging Twitter data and BERT models for depression prediction, owing to the library’s popularity in the NLP community.

Ghosh et al. [3] introduce an Android mobile application that adapts to rapidly changing financial landscapes. The app employs innovative machine learning algorithms for sentiment analysis of financial news, integrating aspect-based sentiment analysis and predictive modelling. It offers insights for round-the-clock planning and predictive eligibility assessment for loans or credit cards. By democratizing financial information and decision-making, the research makes advanced analytics accessible to users of all finance expertise levels. Flores and Antunes [4] examine the transformation of the news ecosystem due to digital media, focusing on news apps and social media platforms. They highlight how these changes present challenges and opportunities for journalism and the public. News apps, favoured by younger audiences for convenience and personalized features, exemplify this shift. Despite media organizations’ financial challenges, mobile app advancements offer new opportunities. The study explores young adults’ perceptions of news apps’ uses and relevance in this evolving media landscape.

Molyneux and Haskell [5] examine how mobile devices have transformed news production and consumption. They describe mobile news consumption as ‘snacking,’ involving brief and frequent interactions. Producers must adapt content for small screens, gaining immediate audience access. Although the impact on learning and mobilization is inconclusive, the context-dependent nature of smartphones’ portability and connectivity opens multiple research pathways. Younger generations, growing up with mobile news, challenge these initial conclusions. Peng et al. [6] explore an automatic news generation and fact-checking system using advanced NLP and deep learning. The system enhances news production efficiency and ensures content authenticity. Key technologies include text generation, information extraction, and knowledge graphs. Experiments validate the system’s effectiveness, and future directions highlight further integration and innovation.

He and Abisado [7] address polysemy and feature extraction issues in text sentiment analysis using a BERT-CNN-BiLSTM-Att hybrid model. The BERT model generates dynamic word vectors, which CNN and BiLSTM process for local and global feature extraction. These features are fused and enhanced using an attention mechanism for sentiment categorization of the Douban movie review dataset. The proposed model outperforms others, including Word2Vec-BiLSTM and BERT-CNN,

showing significant accuracy improvements and demonstrating superior sentiment classification performance. Kumar et al. [8] propose a Cross CNN-LSTM model for sarcasm identification in sentiment analysis. The model leverages convolutional neural networks (CNN) and Long Short-Term Memory (LSTM) networks to process data from social networking and blogging platforms. Using Word2Vec for initial word embedding, the model captures higher-level information through CNN layers and long-term dependencies through LSTM. Including dropout, normalization, and a rectified linear unit enhances accuracy.

Rateb et al. [9] evaluate sentiment analysis methodologies during times of crisis, focusing on tweets related to cryptocurrencies during the Russian-Ukrainian War. They analyze over one million tweets spanning three months and consider Bitcoin, Ethereum, and Binance Coin. Two models, CNN-LSTM and SVM, with GloVe and TF-IDF features, are compared with a pre-trained model, Pysentimento. Pysentimento outperforms the other models in accuracy. Additionally, they employ Google Trends data and cryptocurrency metrics for price prediction analysis using the SARIMA model. Their findings suggest that sentiment analysis using machine learning provides valuable insights for cryptocurrency price forecasting and trading strategies, particularly amidst geopolitical events and market volatility. Susandri et al. [10] introduce a hybrid CNN-BiLSTM model for text sentiment classification using data from WhatsApp groups. They explore five deep learning models and propose a model with enhanced feature extraction and hybrid architecture. Achieving an 88% accuracy on their dataset, their model outperforms previous studies, demonstrating its effectiveness in sentiment analysis.

The studies reviewed demonstrate a wide range of applications for sentiment analysis. Manikandan and Priscila [1] and Rateb et al. [9] illustrate how social media platforms serve as a rich data source for understanding public sentiment during events like the COVID-19 pandemic and geopolitical crises. Their work underscores the critical role of real-time sentiment analysis in crisis management, public health communication, and cryptocurrency forecasting. For instance, Rateb et al. [9]’s evaluation of cryptocurrency sentiment during the Russian-Ukrainian conflict showcases the potential of sentiment analysis to provide predictive insights into financial markets under volatile conditions. Beyond crisis scenarios, sentiment analysis also applies to mental health and finance. Pourkeyvan et al. [2] explore using social media data to predict early symptoms of mental disorders, showcasing how sentiment analysis can assist in mental health diagnostics. Ghosh et al. [3] apply sentiment analysis to financial news, enabling users to make informed decisions about loans and credit. These applications highlight sentiment analysis as a personal and institutional decision-making tool, from health diagnostics to financial risk assessment.

Across these studies, there is a noticeable emphasis on integrating traditional machine learning algorithms with deep learning techniques to improve the accuracy and reliability of sentiment analysis models. Basic models like TextBlob, used by Manikandan and Priscila [1], provide quick sentiment assessments but often lack the depth to handle nuanced contexts such as sarcasm or polysemy. Advanced models, including BERT-based architectures (as used by Pourkeyvan et al. [2] and He and Abisado [7]), demonstrate a marked improvement in accuracy by leveraging pre-trained language models capable of understanding context through bidirectional training. Several studies also introduce hybrid models that combine multiple techniques to address specific challenges. He and Abisado [7] propose a BERT-CNN-BiLSTM-Attention hybrid model to mitigate the impact of polysemy and enhance feature extraction. The integration of CNNs for local feature extraction and BiLSTM networks for long-term dependency management allows these models to capture granular and global information, improving sentiment classification accuracy. Similarly, Kumar et al. [8] develop a Cross CNN-LSTM model specifically for sarcasm detection, highlighting the importance of selecting appropriate models based on data characteristics.

The reviewed literature points to several common challenges in sentiment analysis, many of which are tied to the complexities of natural language. Polysemy, sarcasm, and ambiguous language often hinder the effectiveness of traditional sentiment analysis methods. Researchers have turned to deep learning models that utilize contextual embeddings to address these issues. For example, the BERT models highlighted by Pourkeyvan et al. [2] and He and Abisado [7] demonstrate the capability to disambiguate terms and understand context, reducing misclassification rates. Data quality and feature extraction also present challenges. Susandri et al. [10] emphasize the need for enhanced feature extraction to improve model performance on informal communication platforms like WhatsApp. They demonstrate that hybrid architectures, which fuse various deep learning components, are particularly effective in handling diverse text formats, from casual conversations to formal news articles. Table 1 illustrates the summary of the Literature Review.

Table 1: Summary of Literature Review

Author(s)	Methodology	Challenges
Manikandan and Priscila [1]	Sentiment analysis using Flair PyTorch, emojis, and TextBlob	Limited applicability beyond COVID-19 tweets
Pourkeyvan et al. [2]	Mental health disorder prediction using Hugging Face Transformers	Highly dependent on social media data quality and quantity
Ghosh et al. [3]	Aspect-Based Sentiment Analysis (ABSA) for financial news	Complexity may hinder usability for non-technical users

Flores and Antunes [4]	Study on young adults' perceptions of news apps	Traditional media faces challenges in adapting to digital trends
Molyneux and Haskell [5]	Analysis of Mobile News Consumption Patterns	Unclear impact on learning and civic engagement
Peng et al. [6]	Automatic news generation and fact-checking using NLP	Ethical concerns regarding AI-generated content
He and Abisado [7]	Sentiment analysis using the BERT-CNN-BiLSTM-Att model	High computational cost and resource-intensive
Kumar et al. [8]	Sarcasm detection using a cross CNN-LSTM model	Requires significant processing power and fine-tuning
Rateb et al. [9]	Sentiment analysis for crisis events using SVM, CNN-LSTM, and Psyntimiento	Results fluctuate due to external geopolitical and economic factors
Susandri et al. [10]	Sentiment classification on WhatsApp messages using Hybrid CNN-BiLSTM	Limited generalizability beyond WhatsApp conversations

3. Methodology

The proposed framework is a web application designed to analyze daily news articles and classify them as positive, negative, or neutral in sentiment, providing users with balanced insights. By leveraging advanced Natural Language Processing (NLP) techniques, the project aims to enhance users' understanding of current events, helping mitigate the potential for negativity overload in news consumption. This module description offers an overview of the project's objectives, architecture, and key functionalities, underscoring the technologies that make it effective in sentiment analysis and summarization. It includes several core functions: Sentiment Analysis, which classifies news articles by emotional tone to provide insight into the news landscape; Summarization, which generates brief, impactful summaries to make news accessible quickly; User Interaction, which prioritizes an intuitive interface for easy exploration of news content; and Emotional Well-being, which counters the mental health impacts of negative news by giving a balanced overview of events.

Figure 1 illustrates the Basic Building Blocks of the Goodness model. The proposed architecture consists of multiple integrated modules: The Data Collection Module gathers news articles from RSS feeds and web scraping to ensure a diverse range of news sources. The Data Processing Module includes sentiment analysis and summarization for effective data interpretation. At the same time, the User Interface Module organizes the presentation of news, sentiment data, and summaries in a clear and visually engaging way. The Database Module stores processed articles for easy retrieval and display, enhancing the app's efficiency and scalability. The Data Collection Module gathers news articles from RSS feeds and web scraping using the Python 'feed-parser' and 'trafilatura' libraries. Feed-parser handles RSS parsing to keep the application responsive while maintaining source diversity, and trafilatura extracts key content from each article, eliminating ads and other non-essential information. Each article is structured to store metadata (title, link, summary, full text, sentiment), facilitating efficient storage and retrieval.




Figure 1: Basic block diagram of Goodness framework

Sentiment Analysis is a vital feature of the “Goodness” project, achieved using a dual approach: TextBlob for initial sentiment categorization and Logistic Regression for deeper analysis. TextBlob uses polarity scores to classify sentiment quickly into positive, negative, or neutral categories, while Logistic Regression, implemented through 'sklearn,' further analyzes articles with more nuanced content. Logistic regression, trained on datasets like the NLTK movie reviews, allows the system to handle subtle emotional distinctions within articles, increasing classification accuracy. For summarization, the project leverages the BART model (Bidirectional and Auto-Regressive Transformers), which is known for generating effective text summaries. BART processes each article by tokenizing it and producing a summary that retains essential points in an accessible format.

The model, loaded from pre-trained versions, operates efficiently without extensive retraining, generating summaries that condense articles while preserving core context. The User Interface Module is central to user engagement. Built with HTML and CSS for a responsive design, the interface is visually engaging, displaying headlines, summaries, and sentiment classifications with colour-coded indicators—green for positive, yellow for neutral, and red for negative. Dynamic content and interactive animations add to the experience, encouraging users to explore news easily and providing feedback through notifications for newly fetched articles.

Database Module stores and manages news articles and related data for quick access and retrieval. While optional in the early stages, a formal database can support future scalability, allowing for features like personalized feeds and analytics. As the app grows, the database will facilitate efficient storage, enabling rapid content updates and making it suitable for future expansion. Deployment involves cloud-based hosting solutions, such as AWS or Heroku, to ensure reliability and user accessibility without relying on local servers. Ngrok creates secure tunnels for real-time testing and webhook verification during development, aiding a smooth deployment process. This proposed work offers a comprehensive approach to news sentiment analysis, enhancing user understanding of emotional tones while providing concise news summaries. By integrating NLP with a user-friendly interface, the application promotes a more balanced news consumption experience, encouraging users to engage with diverse content. The modular architecture, real-time data handling, and focus on user mental well-being contribute to a healthier approach to news reading and make “Goodness” a valuable tool in mitigating the effects of negativity in media.




Figure 2: Functional flow diagram

Figure 2 depicts the functional flow diagram of the Goodness model. The functional flow diagram provides a detailed workflow of the Goodness system, an advanced web-based tool for news summarization and sentiment analysis. The system is structured into three major phases: data preparation and model setup, core processing, and deployment & web interaction. The first phase ensures that all necessary libraries, datasets, and models are properly initialized. This begins with importing essential Natural Language Processing (NLP), machine learning, and web development libraries to facilitate text processing and analysis. Next, Google Drive is mounted, allowing access to large datasets and pre-trained models stored in the cloud. The BART model and tokenizer are then loaded, which are crucial for text summarization tasks. Additionally, NLTK datasets are downloaded to provide linguistic resources for sentiment classification. Finally, the collected data is prepared and split into training and testing sets, ensuring the system is trained efficiently and validated for accuracy before further processing.

The system implements machine learning models to analyze and summarize news articles in the core processing phase. First, a text vectorizer transforms raw textual data into numerical features, making it suitable for machine learning algorithms. This vectorized data is then used to train a Logistic Regression model, which plays a crucial role in sentiment classification by determining whether the summarized text conveys a positive, negative, or neutral sentiment. Following this, a Flask-based web application is initialized, providing users with an interactive platform to input and process news articles dynamically. The system also defines two key functions: a summarization function that utilizes the BART model to extract the most relevant parts of a news article and a sentiment analysis function that applies the trained logistic regression model to the summarized content. This two-step approach ensures that the system reduces the article’s length while maintaining its essence and accurately classifies its emotional tone, allowing for an in-depth understanding of news sentiment.

The final phase focuses on real-time news integration, deployment, and accessibility. The system is designed to fetch live news articles from an RSS feed, making it adaptable to evolving global events. Once the articles are retrieved, they undergo summarization and sentiment analysis, providing users with a concise yet sentiment-rich understanding of the news. These analyzed articles are stored in a global variable, ensuring efficient retrieval and minimizing redundant processing. A web interface is built using Flask to enhance usability, offering an intuitive user experience. Finally, the system is deployed using ngrok, enabling remote access to the web application from anywhere. This end-to-end approach ensures that users can quickly

digest the latest news with sentiment-aware insights, making informed decisions without reading lengthy articles. The structured pipeline of the Goodness system combines NLP, machine learning, and web technologies, providing an automated, real-time, and interactive news analysis solution that is both efficient and user-friendly.

The user interface, developed using Flask, engages in its clean, responsive layout, organizing summaries and their sentiment classification. The sentiments are colour-coded- green for positive, yellow for neutral, and red for negative- to quickly understand their stories’ emotional context. It looks visually attractive and user-friendly due to smooth animations and intuitive navigation features. Precision, recall, and F1-score are evaluation metrics that will ensure “Goodness” is effective in measuring the performance of the sentiment analysis model. User feedback would constitute an intrinsic part of the process, with iterative improvements in functionality and accuracy subsequently happening. Distinctly different from existing solutions, this system offers a holistic news consumption experience. It can pull together real-time updates while enhancing readability and nuanced sentiment analysis for a comprehensive yet simplified view of news articles. “Goodness” aims to create informed decision-making and regular engagement through a thoughtfully imagined user interface.

4. Results and Discussions

TextBlob is a lightweight and user-friendly natural language processing (NLP) tool that excels in sentiment analysis due to its high recall. It effectively captures a wide range of positive sentiments with minimal omission. This characteristic makes it particularly useful for applications where identifying and filtering neutral emotions is crucial. However, its lower precision indicates a tendency to overpredict positive sentiments, which may lead to inaccuracies when distinguishing between genuinely positive and neutral or mixed sentiments.

Table 2: Performance metrics of TextBlob Classifier

Model	Accuracy	Precision	Recall	F1 Score
TextBlob	0.6175	0.569767	0.975124	0.719266

Despite this limitation, TextBlob remains a valuable tool for quick and efficient sentiment assessment, particularly when used with more sophisticated models like Logistic Regression. By leveraging its strengths in the recall, TextBlob ensures comprehensive sentiment detection while adding a secondary, more precise model mitigates its tendency to overestimate positivity. This hybrid approach balances accuracy and efficiency, making it a practical choice for real-time sentiment analysis tasks, such as news article summarization and classification. Table 2 lists the performance metrics of TextBlob Classifier.




Figure 3: TextBlob Confusion Matrix

Figure 3 depicts the confusion matrix of TextBlob. The TextBlob demonstrates a strong ability to classify positive sentiments, correctly identifying 196 positive instances. However, its major weakness is its tendency to misclassify negative sentiments as positive, with 148 false positives. This indicates a significant bias toward positive sentiment predictions. The model’s high recall for positive sentiment is beneficial, but its very low precision for negative sentiment makes it unreliable for detecting negative sentiments accurately.

Table 3: Performance metrics of NRCLex Classifier

Model	Accuracy	Precision	Recall	F1 Score
NRCLex	0.6425	0.59864	0.87562	0.71111

Table 3 lists the performance metrics of the NRCLex Classifier. The NRCLex is a sentiment analysis tool that provides a balanced approach, offering better accuracy and precision than TextBlob while maintaining a moderate recall rate. Unlike TextBlob, which tends to overpredict positive sentiments, NRCLex’s slightly lower recall reduces the likelihood of misclassifying neutral or mixed sentiments as overly positive. This makes it a more reliable choice for general sentiment detection, where minimizing false positives is important. However, while NRCLex performs consistently across different datasets, it does not particularly excel in any specific aspect compared to more advanced deep learning-based sentiment models. Its strength lies in serving as a stable baseline for sentiment analysis, making it a practical option for applications that require quick and straightforward sentiment classification without the computational complexity of neural network-based models. Nonetheless, its lack of depth in capturing nuanced emotions limits its effectiveness in domains where fine-grained sentiment distinctions are necessary, such as detecting sarcasm or complex emotional tones in subjective text.


**Figure 4:** NRCLex Confusion Matrix

Figure 4 demonstrates the NRCLex Confusion Matrix. NRCLex performs better than TextBlob in recognizing negative sentiments, correctly classifying 81 negative instances compared to TextBlob’s 51. However, it still exhibits a high false positive rate, with 118 negative cases incorrectly classified as positive. While this model balances predictions slightly better than TextBlob, it favours positive sentiment too much, leading to misclassifications of negative sentiments.

Table 4: Performance Metrics of Logistic Regression Classifier

Model	Accuracy	Precision	Recall	F1 Score
Logistic Regression	0.83	0.84103	0.81592	0.82828

Table 4 lists the performance metrics of the Logistic Regression classifier. Logistic regression is a widely used machine learning algorithm for sentiment analysis because it provides a balanced approach with high accuracy, precision, and recall. Unlike rule-based approaches like TextBlob, which may overpredict positive sentiments, Logistic Regression effectively captures sentiment trends while minimizing false predictions. Its strength lies in its ability to generalize well across different datasets, ensuring consistent performance. When paired with TextBlob, it compensates for the latter’s tendency to overpredict by refining sentiment classification, particularly for neutral sentiments. This hybrid approach enhances sentiment analysis by combining the linguistic simplicity of TextBlob with the robust statistical modelling of Logistic Regression. As a result, this combination provides a nuanced understanding of sentiment distribution, making it particularly useful in applications where distinguishing between subtle emotional tones is crucial. Logistic regression’s computational efficiency also ensures scalability, making it suitable for real-time sentiment detection tasks in large-scale social media monitoring and news sentiment analysis applications.




Figure 5: Logistic Regression Confusion Matrix

Figure 5 shows the confusion matrix of the Logistic Regression model. Logistic regression provides one of the best-balanced sentiment classification results. It significantly reduces false positives and false negatives, correctly identifying 168 negative and 164 positive instances. However, it still misclassifies 37 positive and 31 negative instances as positive. Despite these misclassifications, this model achieves a strong trade-off between precision and recall, making it one of the most reliable sentiment classification models.

Table 5: Performance metrics of Multinomial Naïve Bayes Classifier

Model	Accuracy	Precision	Recall	F1 Score
Multinomial Naïve Bayes	0.815	0.83957	0.7811	0.80928

Table 5 lists the performance metrics of the Multinomial Naïve Bayes classifier. Multinomial Naïve Bayes is a probabilistic classifier that performs well in sentiment analysis by leveraging word frequency distributions. It exhibits strong precision, accurately predicting positive sentiments with minimal false positives. However, its recall is slightly lower, indicating some positive sentiments may be missed. This trade-off makes it particularly useful in scenarios where precision is more critical than recall, such as applications that require highly accurate sentiment classification rather than capturing the full range of sentiments. The model's simplicity and efficiency make it well-suited for large-scale text classification tasks, including spam filtering and topic categorization. However, its reliance on word frequency without considering contextual meaning can limit its effectiveness for nuanced sentiment detection. Despite this, when combined with other models, such as logistic regression or deep learning approaches, Multinomial Naïve Bayes can enhance overall sentiment analysis performance by providing a reliable baseline for classification.




Figure 6: Multinomial Naïve Bayes Confusion Matrix

Figure 6 depicts the confusion matrix of Multinomial Naïve Bayes. Multinomial Naïve Bayes performs similarly to Logistic Regression, with slightly better negative sentiment classification, correctly identifying 169 negative instances. However, it exhibits a slightly higher false negative rate, misclassifying 44 positive instances as negative. This suggests that the model is conservative in predicting positive sentiment, meaning it is more likely to label a positive sentiment incorrectly as negative.

Table 6: Performance metrics of SVM Classifier

Model	Accuracy	Precision	Recall	F1 Score
SVM	0.735	0.77778	0.66169	0.71505

Table 6 shows the performance metrics of the SVM Classifier. SVM effectively differentiates between sentiment classes by finding an optimal decision boundary, ensuring minimal false positives. However, its recall is lower, meaning it tends to be conservative in its predictions and may fail to capture some sentiments, particularly nuanced or ambiguous ones. This characteristic makes SVM ideal for applications that prioritize reducing false positives, such as detecting critical sentiments in financial news or medical reports. Despite its advantages, SVM is computationally expensive, requiring significant processing power and time for large datasets. It also lacks probabilistic output, making it less interpretable than logistic regression models. While SVM can be a reliable sentiment classifier, its recall and computational complexity limitations often make logistic regression or deep learning-based approaches more favourable in real-world sentiment analysis tasks.




Figure 7: SVM Confusion Matrix

Figure 7 depicts the confusion matrix of SVM. SVM excels in identifying negative sentiments, correctly classifying 161 negative cases. However, its major weakness is its high false negative rate, misclassifying 68 positive instances as negative. This suggests that the model is overly conservative, frequently underestimating positive sentiment. While it effectively identifies negativity, it struggles to balance its predictions, making it less suitable for cases where positive sentiment detection is equally important.

Table 7: Performance metrics of Linear SVC Classifier

Model	Accuracy	Precision	Recall	F1 Score
Linear SVC	0.82	0.82412	0.81592	0.82

Table 7 shows the performance metrics of the Linear SVC Classifier. Linear SVC is a strong performer in sentiment analysis, providing consistent accuracy, precision, and recall across different datasets. Its ability to handle high-dimensional data efficiently makes it a reliable choice for text classification tasks. Unlike standard SVM, Linear SVC is optimized for large-scale problems, making it computationally more efficient. However, it lacks the interpretability of Logistic Regression, as it does not provide probabilistic outputs, which can make sentiment classification less intuitive. While its balanced performance makes it comparable to Logistic Regression, the latter remains preferable when simplicity and interpretability are prioritized. When combined with models like TextBlob, Logistic Regression offers better transparency in sentiment classification, making it a more practical choice for applications requiring explainability.




Figure 8: Linear SVC Confusion Matrix

Figure 8 shows the Linear SVC Confusion Matrix. Linear SVC exhibits a classification pattern similar to standard SVM, effectively classifying negative sentiments while showing many false negatives. This means that many positive sentiments are incorrectly classified as negative. While the model is useful in detecting negative sentiment, its tendency to underestimate positive sentiment makes it less effective in a balanced sentiment analysis task.

Table 8: Performance metrics of DistilBERT Classifier

Model	Accuracy	Precision	Recall	F1 Score
DistilBERT	0.745	0.76191	0.71642	0.73846

Table 8 shows the performance metrics of the DistilBERT Classifier. DistilBERT leverages transformer-based embeddings to achieve a deeper contextual understanding of a text, making it particularly effective for sentiment analysis tasks that require nuanced interpretation. It offers good precision and can accurately identify sentiments without excessive false positives. However, its accuracy is only moderate compared to more optimized models, and its computational requirements are significantly higher. This makes it less ideal for lightweight applications where speed and efficiency are priorities. While DistilBERT excels in complex sentiment classification, it may not always justify the resource cost compared to simpler models like Logistic Regression, which, combined with TextBlob, can efficiently classify neutral sentiments while maintaining interpretability and ease of deployment.




Figure 9: DistilBERT Confusion Matrix

Figure 9 shows the confusion matrix of DistilBERT. DistilBERT Confusion Matrix, a transformer-based model, shows slightly weaker performance than expected, possibly due to insufficient fine-tuning. It correctly classifies 154 negative instances (True Negatives) but has 45 false positive errors. It also correctly predicts 144 positive instances (True Positives) while making 57 false negative errors. While it is still competitive, its error rate suggests it may require additional training data or parameter tuning to match the performance of simpler models like Logistic Regression or SVC. Among the sentiment analysis models tested, Logistic Regression and Multinomial Naïve Bayes exhibit the most balanced performance, effectively classifying positive and negative sentiments while maintaining low misclassification rates. TextBlob and NRCLex show a strong bias toward positive sentiment, frequently misclassifying negative instances, making them less reliable for detecting negativity. SVM and Linear SVC, on the other hand, demonstrate strong negative sentiment classification but suffer from a high false negative rate, often mislabeling positive sentiments as negative. Logistic regression is the most reliable model, providing a good trade-off between precision and recall. Multinomial Naïve Bayes also performs well, slightly favouring negative sentiment classification. However, TextBlob and NRCLex are not well-suited for balanced sentiment analysis due to their tendency to misclassify negative sentiments as positive.

The classification reports in the Goodness project provide a detailed evaluation of the performance metrics, including precision, recall, F1-score, and support for various sentiment analysis models. These reports give insights into how well each model performs across different sentiment categories, clearly understanding their strengths and weaknesses. The TextBlob classification report provides an in-depth analysis of the model's accuracy in categorizing news articles as positive, neutral, or negative, detailing precision (how many selected items are relevant), recall (how many relevant items are selected), and the F1-score, which balances both. Similarly, the NRCLex classification report evaluates its ability to assess emotional tones, offering insights into how accurately it detects specific emotions and the balance between its precision and recall. The logistic regression classification report shows its effectiveness in binary sentiment classification by quantifying its performance through precision, recall, and F1-scores, highlighting areas where it may over or underpredict sentiment categories.

The Multinomial Naïve Bayes classification report provides similar metrics, showcasing its probabilistic approach to sentiment classification and highlighting its balance between correctly identifying sentiments and potential misclassifications. The Support Vector Machine (SVM) classification report delves into the model's ability to differentiate between positive, neutral, and negative sentiments with a particular emphasis on precision and recall, especially in handling ambiguous or borderline cases. The Linear SVC classification report offers a focused analysis of the model's performance, detailing how it balances precision and recall in sentiment prediction. Finally, the DistilBERT classification report provides a high-level overview of its capabilities, demonstrating its advanced handling of sentiment classification with high precision and F1 scores, emphasizing its overall effectiveness. Together, these classification reports offer a thorough comparison of the models, highlighting their predictive power and reliability across various sentiment categories.

Table 9: Comparison of Models based on performance

Model	Accuracy	Precision	Recall	F1 Score
TextBlob	0.62	0.57	0.98	0.72
NRCLex	0.64	0.6	0.88	0.71
Logistic Regression	0.83	0.84	0.82	0.83
Multinomial Naïve Bayes	0.81	0.84	0.78	0.81
Support Vector Machine (SVM)	0.73	0.78	0.66	0.72
Linear SVC	0.82	0.82	0.82	0.82
DistilBERT	0.74	0.76	0.72	0.74

Table 9 shows the comparison of models based on performance. TextBlob achieves a moderate accuracy of 61.75%, making it a simple yet effective tool for sentiment analysis. Its major strength lies in its high recall of 97.51%, which ensures that almost all positive sentiments are detected. However, its low precision of 56.97% leads to frequent false positives, making it less reliable for applications where accuracy in positive classification is crucial. NRCLex slightly improves upon TextBlob, offering an accuracy of 64.25%. It maintains a better balance between precision and recall, resulting in a stable F1 score of 71.11%. Despite this improvement, it still struggles with negative samples, leading to relatively low true negative rates, which may affect overall sentiment classification effectiveness. Logistic regression stands out with a high accuracy of 83%, demonstrating consistent performance across sentiment categories. It effectively identifies positive and negative sentiments with a precision of 84.10% and a recall of 81.59%. However, a slight drop in recall means it occasionally produces false negatives, missing some sentiment classifications.




Figure 10: Model Accuracy Comparison Graph

Multinomial Naïve Bayes delivers a strong accuracy of 81.5%, offering high precision at 83.95% and recall at 78.11%. These metrics ensure reliable predictions across different sentiment categories. However, its recall for positive instances is slightly lower, which reduces its overall F1 score and makes it less effective in capturing all positive sentiments. Support Vector Machine (SVM) achieves a moderate accuracy of 73.5%, excelling in precision at 77.78%. This suggests that it accurately detects positive sentiments while minimizing false positives. However, its recall of 66.17% is lower, leading to the omission of some positive instances, which could limit its reliability in comprehensive sentiment analysis. Linear SVC performs well with an accuracy of 82%, showing strong and balanced precision and recall at 82.41% and 81.59%, respectively. This robustness ensures a reliable sentiment classification across different datasets. However, its marginal difference from Logistic Regression suggests that it may not provide significant additional benefits over simpler models. Figure 10 illustrates the comparison of the Accuracy graph.

DistilBERT, leveraging transformer-based embeddings, achieves a decent accuracy of 74.5%. It maintains a balanced precision and recall, making it stable for sentiment classification tasks. However, its recall is lower than models like Logistic Regression and Linear SVC, negatively impacting its F1-score. Additionally, its computational demands may make it less suitable for resource-constrained environments. The sentiment analysis models exhibit varying strengths and weaknesses. TextBlob prioritizes recall, capturing nearly all positive sentiments but suffering from frequent false positives. NRCLex offers slightly better accuracy with improved balance, though it struggles with negative samples. Logistic Regression and Linear SVC provide the highest accuracy (~82-83%) with balanced precision and recall, making them reliable choices for sentiment classification. Multinomial Naïve Bayes performs well with strong precision but slightly lower recall. SVM demonstrates high precision but lower recall, missing some positive sentiments.

DistilBERT, while leveraging deep learning, shows moderate performance and higher computational demands, making it less efficient than simpler models. In General, Logistic Regression and Linear SVC offer the best balance of precision, recall, and computational efficiency. TextBlob and NRCLex are useful for quick, broad sentiment detection but lack robustness for detailed classification. Multinomial Naïve Bayes and SVM are effective in specific scenarios where precision is more critical than recall. While DistilBERT provides a deeper contextual understanding, its performance does not significantly surpass simpler models, making it less practical for real-time applications. Ultimately, the choice of model depends on the specific needs—whether prioritizing precision, recall, or computational efficiency. Table 10 lists a comparison of models based on Confusion Matrix Value.

Table 10: Comparison of Models based on Confusion Matrix Value

Model	True Positives (TP)	False Positives (FP)	True Negatives (TN)	False Negatives (FN)
TextBlob	196	148	51	5
NRCLex	176	118	81	25
Logistic Regression	164	31	168	37

Multinomial Naïve Bayes	157	30	169	44
Support Vector Machine (SVM)	133	38	161	68
Linear SVC	164	35	164	37
DistilBERT	144	45	154	57

The comparison of models in the Goodness project provides an insightful evaluation of various sentiment analysis models, including TextBlob, NRCLEX, Logistic Regression, Multinomial Naïve Bayes, SVM, Linear SVC, and DistilBERT. Each model is assessed using key metrics like precision, recall, and F1-score. Simpler models like TextBlob and NRCLEX offer basic sentiment classification, while more advanced models like DistilBERT outperform them with higher accuracy and precision in complex scenarios. Logistic Regression and Naïve Bayes show consistent performance, and SVM models excel at handling boundary cases, making the comparison crucial for choosing the optimal model. Figure 11 lists the First Page colour-coded with the Latest News Headlines and Summary.



Figure 11: First Page – Colour-Coded Latest News Headlines and Summary

This image showcases the application’s first page, presenting users with the latest news headlines and summaries. Articles are colour-coded based on sentiment analysis—green for positive, yellow for neutral, and red for negative news. The concise summaries allow users to quickly assess the content and decide whether to read the full article. The colour-coded design helps users filter and focus on uplifting news, aligning with the project’s goal of providing positive content (Figure 12).




Figure 12: Neutral News – Full Article

The Neutral News page displays a full news article classified as neutral by the sentiment analysis module. The article has its headline, publication date, and content in a clean, easy-to-read format. The colour coding remains yellow, indicating neutral sentiment (Figure 13).




Figure 13: Negative News – Full Article

This image displays a full news article classified as negative, highlighted in red. Despite the Goodness project's focus on positive content, negative articles can still be accessed. The article page features the headline, publication date, and full content. Negative articles, while available, are minimized to maintain the application's uplifting focus. The red coding ensures the user knows the article's negative sentiment (Figure 14).




Figure 14: Positive News – Full Article

The Positive News page presents a full article classified as positive by the sentiment analysis system, colour-coded green. This article displays its headline, publication date, and content. The page promotes an encouraging and optimistic reading experience, aligning with the project's mission to provide positive news. Users can easily navigate and enjoy uplifting stories, which form the application's core.

5. Conclusion

The proposed Goodness framework combines natural language processing and sentiment analysis to create a concise, sentiment-aware news platform that addresses information overload challenges and emotionally charged content's psychological impact. By using AI models like BART for text summarization and a hybrid approach with TextBlob and Logistic Regression for sentiment analysis, "Goodness" provides users with contextually accurate summaries and sentiment tags (positive, neutral, or negative), helping them stay informed while managing emotional exposure. The user interface, developed with Flask, presents these summaries and sentiment tags in a visually engaging, colour-coded format, making it easy for readers to understand the emotional tone of the news. With the potential for future improvements like expanded emotional classifications, multilingual support, personalized recommendations, and real-time model training based on user feedback, "Goodness" lays a strong foundation for a mindful, emotionally aware news experience, contributing to reader well-being in an increasingly information-saturated world.

Acknowledgment: N/A

Data Availability Statement: The data for this study can be made available upon request to the corresponding author.

Funding Statement: This manuscript and research paper were prepared without any financial support or funding.

Conflicts of Interest Statement: The authors have no conflicts of interest to declare.

Ethics and Consent Statement: This research adheres to ethical guidelines, obtaining informed consent from all participants.

References

1. N. Manikandan and S. S. Priscila, "Sentiment Analysis of COVID-19 Tweets Through Flair PyTorch, Emojis, and TextBlob," in *Explainable AI Applications for Human Behavior Analysis*, pp. 275–289, IGI Global, Hershey, Pennsylvania, United States of America, 2024.
2. A. Pourkeyvan, R. Safa, and A. Sorourkhah, "Harnessing the power of Hugging Face transformers for predicting mental health disorders in social networks," *IEEE Access*, vol. 12, no. 2, pp. 28025–28035, 2024.
3. S. Ghosh, M. Sengupta, and S. Chakrabarti, "Aspect-based sentiment analysis Android app for financial news," *Journal of Research Administration*, vol. 6, no. 1, p. 3, 2024.
4. A. M. M. Flores and I. E. Antunes, "News in The Pocket: a Study Of Young Adult's Perceptions Of The Uses And Relevance Of News Apps," *Inteligência Artificial e Jornalismo Móvel*, LabCom, London, England, United Kingdom, 2024.
5. L. Molyneux and A. Haskell, "News on Smartphones: Exposure, Consumption, and Information through Mobile Media," in *The Mobile Media Debate*, pp. 25–36, Routledge, London, England, United Kingdom, 2024.
6. X. Peng, Q. Xu, Z. Feng, H. Zhao, L. Tan, Y. Zhou, Z. Zhang, C. Gong, and Y. Zheng, "Automatic news generation and fact-checking system based on language processing," vol. 1, no. 5, pp. 1–10, 2024.
7. A. He and M. Abisado, "Text sentiment analysis of Douban film short comments based on BERT-CNN-BiLSTM-Att model," *IEEE Access*, vol. 12, no. 3, pp. 45229–45237, 2024.
8. S. Kumar, A. K. Singh, S. Bhushan, and V. K. Singh, "A cross CNN-LSTM model for sarcasm identification in sentiment analysis," in *Artificial Intelligence, Blockchain, Computing and Security*, pp. 322–328, CRC Press, London, United Kingdom, 2023.
9. M. N. Rateb, S. Alansary, M. K. Elzouka, and M. Galal, "Sentiment Analysis in Times of Crisis: Evaluating SVM, CNN-LSTM, and Psyntimiento," vol.2, no.2, pp.1–28, 2024.
10. S. Susandri, S. Defit, and M. Tajuddin, "Enhancing text sentiment classification with hybrid CNN-BiLSTM model on WhatsApp group," *J. Adv. Inf. Technol.*, vol. 15, no. 3, pp. 355–363, 2024.
11. "Bad News Bias and Its Impact," *Vox*, 2023. [Online]. Available: <https://www.vox.com/the-highlight/23596969/bad-news-negativity-bias-media>. [Accessed by 11/03/2024]
12. "Negativity in News Consumption," *Nature*, 2023. [Online]. Available: <https://www.nature.com/articles/s41562-023-01538-4>. [Accessed by 11/03/2024]
13. C. Higson, "What is Negativity Bias and How Can We Overcome It?," *Consult Higson*, Apr. 3, 2023. [Online]. Available: <https://consulthigson.com/2023/04/03/what-is-negativity-bias-and-how-can-we-overcome-it/>. [Accessed by 11/03/2024]
14. Elgwaro, "Kenyan Media Outlets Must Avoid Negative News to Save Our Country," *Medium*, 2023. [Online]. Available: <https://medium.com/@elgwaro/kenyan-media-outlets-must-avoid-negative-news-to-save-our-country-879be82e187c>. [Accessed by 11/03/2024]
15. T. Palm, "Negative or Positive Headlines: Which Do You Choose?," *LinkedIn*, 2023. [Online]. Available: <https://www.linkedin.com/pulse/negative-positive-headlines-which-do-you-choose-tony-palm->. [Accessed by 11/03/2024]