



Sentiment Analysis of the 2022 Fuel Price Hike Using the Naïve Bayes Classifier

Alia Lestari*, Muhammad Hajarul Aswad, and Subekti Masri

Universitas Islam Negeri Palopo, Indonesia

Abstract

This study examines public opinion towards Indonesia's 2022 fuel price hike using social media analytics and assesses the performance of a supervised machine learning classifier for policy-oriented sentiment analysis. The research aimed to answer the following two questions (1) What was the prevailing public sentiment articulated on Twitter after fuel pricing announcement? and (2) How well is a Naïve Bayes classifier able to classify sentiment polarity in this domain? We employed a quantitative cross-sectional design utilizing Twitter data obtained from 3–4 September 2022 via the hashtags #hargabbm and #bbmnaik (#fuelprice, #fuelpricesincrease). Ultimately, after preprocessing, there were 1,867 unique tweets out of the 2,003 retrieved ones. Training data consisted of a total of 489 manually labeled tweets, and 1,378 for testing. Tokenization and TF-IDF weighting were performed on text data, while the sentiment classification was done using a Gaussian Naïve Bayes model and evaluated through confusion matrix metrics. The results suggest that public sentiment was overwhelmingly negative during the analysis period and that the classifier reached an accuracy of 94.89% with a precision of 73.40%, recall of 100%, and F1-score of 84.66%. These findings show that probabilistic text classification offers evidence about whether the public unanimously supports economically sensitive policies (or not), with intensity and salience meaningfully specified.

Keywords: Energy policy; Naïve Bayes classifier; Sentiment analysis; TF-IDF; Twitter.

Copyright © 2025 by Authors, Published by CAUCHY Group. This is an open access article under the CC BY-SA License (<https://creativecommons.org/licenses/by-sa/4.0>)

1. Introduction

Energy pricing policy remains one of the most politically sensitive and economically consequential policy domains worldwide. Fuel price adjustments directly influence transportation costs, inflation rates, production expenses, and household purchasing power, thereby affecting both macroeconomic stability and everyday life [1], [2]. Globally, governments face persistent trade-offs between fiscal sustainability, subsidy reform, and social acceptance amid volatile international oil market [3], [4]. In developing countries, these tensions are often more pronounced because fuel subsidies function not only as economic instruments but also as mechanisms of social protection [5]. In Indonesia, periodic fuel price increases have repeatedly triggered public debate and protest, underscoring the broader issue of policy legitimacy in economically sensitive reforms [6], [7].

In this context, sentiment analysis refers to the computational identification and classification of opinions or attitudes expressed in textual data into categories such as positive or negative [8].

*Corresponding author. E-mail: alia_lestari@uinpalopo.ac.id

When applied to public policy, sentiment analysis enables large-scale measurement of societal reactions to governmental decisions [9], [10]. For example, a tweet stating that “fuel price hikes burden ordinary citizens” reflects negative sentiment, whereas a tweet arguing that “subsidy reform is necessary for fiscal health” reflects positive sentiment. These individual expressions, when aggregated, provide a structured representation of public response to policy change.

A growing body of literature has examined sentiment in energy-related contexts. In one study, Nguyen investigated whether social media sentiment influences energy market volatility using hybrid sentiment scoring and econometric modeling [11]. The findings showed that negative sentiment was associated with increased price instability, implying that public discourse may affect economic dynamics. Similarly, Deniz and Stengos analyzed institutional communication sentiment and its asymmetric impact on electricity prices, concluding that negative narratives amplify price reactions [12]. These studies suggest that sentiment carries measurable economic consequences.

Beyond market effects, other scholars have explored sentiment as an indicator of public trust and policy evaluation. In one study, Priadana and Rizal analyzed Twitter sentiment toward government performance during the COVID-19 pandemic using supervised machine learning and reported that online sentiment corresponded with public satisfaction trends [13]. Their findings imply that digital discourse can function as a proxy for public legitimacy assessment. Methodologically, Agarwal and Mittal compared Naïve Bayes and Support Vector Machines for text classification and demonstrated that probabilistic classifiers remain effective in high-dimensional textual data [14]. Similarly, Patel and Chhinkaniwala found that Naïve Bayes achieved competitive performance due to computational efficiency and robustness to sparse features [15].

Taken together, prior research establishes three key insights. First, sentiment expressed in digital platforms reflects meaningful economic and social dynamics. Second, machine learning methods, particularly supervised classifiers, can reliably categorize textual polarity. Third, energy-related discourse often exhibits heightened emotional intensity, particularly in contexts involving price adjustments.

However, important gaps remain. Most energy-sentiment studies focus on developed economies or financial market outcomes rather than direct public reaction to specific government policy announcements in developing countries. Moreover, many studies rely on extended temporal datasets, potentially obscuring the immediacy of public reaction during the critical early phase following a policy announcement. In Indonesia, empirical evidence on real-time sentiment toward fuel price reform using supervised machine learning remains limited. This gap is important because early-stage public reaction may shape policy communication strategies, political legitimacy, and long-term acceptance.

Accordingly, this study aims to analyze public sentiment expressed on Twitter in response to Indonesia’s 2022 fuel price increase and to evaluate the statistical effectiveness of a Naïve Bayes classifier in identifying sentiment polarity within this domain. The study addresses two research questions: (i) What is the dominant public sentiment expressed on Twitter following the fuel price announcement? and (ii) How effectively can a Naïve Bayes classifier classify sentiment in this policy context? Based on the economic sensitivity of fuel pricing, it is hypothesized that negative sentiment will dominate public discourse and that the classifier will demonstrate high predictive performance.

To address these questions, a quantitative cross-sectional design is employed. Twitter data are collected during the immediate post-announcement period and subjected to systematic preprocessing and TF-IDF feature extraction. A Gaussian Naïve Bayes classifier is trained on manually labeled data and evaluated using standard classification metrics. It is predicted that probabilistic modeling will yield strong accuracy while capturing the dominant polarity pattern.

This study contributes both theoretically and practically. Theoretically, it integrates applied statistical modeling with public policy analysis in a developing-country energy context.

Methodologically, it demonstrates the continued relevance of generative probabilistic classifiers for real-time policy sentiment analysis. Practically, it offers policymakers an evidence-based approach to monitoring public reaction, thereby supporting more responsive and transparent governance in economically sensitive reforms.

2. Methods

This section describes the methodological framework adopted to address the research questions. It outlines the research design, data collection procedures, preprocessing techniques, feature extraction strategy, classification modeling, and evaluation metrics. Each stage is presented sequentially to ensure analytical transparency and reproducibility of the sentiment analysis process.

2.1. Research Design

This study employs a quantitative approach, using sentiment analysis, to map public opinion on Twitter regarding the 2022 fuel price increase policy. The first stage is data crawling. Data were collected via Google Colab using the #hargabbm and #bbmnaik keywords during September 3–4, 2022, yielding 2,003 relevant tweets. The second stage is data preprocessing, which involves cleaning and removing duplicates. In data cleaning, the data is cleaned of links, URLs, mentions, symbols, emojis/emoticons, and excess spaces. When eliminating duplicates, the same sentence/tweet will be deleted, leaving only one unique instance. The third stage is data labeling.

At this stage, the data will be divided into two groups: training and testing, with positive or negative categories. Tweets with positive categories are interpreted as support or acceptance of the policy, while negative categories indicate rejection or criticism. The labeling of the training data is done manually, while the rest, hereinafter referred to as data testing, is performed automatically using a model trained on the training data with the Naïve Bayes Classifier algorithm. Previously, data extraction was performed using tokenization, yielding datasets containing several important attributes or words. The determination of this critical word utilizes the Term Frequency–Inverse Document Frequency (TF–IDF) rule, which calculates the probability of a word appearing in a single tweet and across the entire corpus of tweets. The last stage is the sentiment analysis stage. At this stage, the performance data, consisting of the levels of accuracy, precision, recall, and F1-score, will be presented. The entire methodological sequence is summarized in the flowchart presented in [Fig. 1](#).

2.2. Data Crawling

The dataset comprised public tweets retrieved from the Twitter platform. Data crawling was conducted using Google Colab as the computational environment. Tweets were collected using the hashtags #hargabbm and #bbmnaik (#fuelprice and #fuelpricesincrease), which were selected to capture discussions directly related to the fuel price increase policy. The data collection period was restricted to 3–4 September 2022 to ensure temporal consistency in public discourse following the policy announcement.

Tweet retrieval was performed by executing a crawling script within Google Colab that queried Twitter using the specified hashtags and date filters. Only tweets containing textual content were retained. Retweets without additional commentary were excluded to avoid duplication of identical content. The crawling process resulted in 2,003 relevant tweets, which were exported and stored in tabular format for subsequent processing.

2.3. Text Preprocessing

All preprocessing procedures were conducted using RapidMiner Studio Version 10.1. The preprocessing stage aimed to standardize textual data and remove noise prior to feature extraction




Fig. 1: Research Flowchart

and classification.

Text cleaning was first applied to remove non-informative components. All URLs or hyperlinks embedded within tweets were removed. User mentions beginning with the “@” symbol were deleted. Non-alphabetic symbols and punctuation marks were eliminated. Emoji and emoticon characters were removed to ensure uniform textual representation. Excessive whitespace was normalized by converting multiple spaces into a single space.

After cleaning, duplicate removal was performed. Tweets with identical textual content were identified through exact string matching. When duplicate tweets were detected, only one instance was retained while all other identical entries were deleted. This step ensured that each observation in the dataset represented a unique textual record and prevented bias arising from repeated content. The following flowchart provides a clear overview and simplifies understanding of the sequence of stages in the research process.

2.4. Data Labelling

Following preprocessing, the dataset was prepared for supervised classification. The data were divided into two subsets: training data and testing data. The training dataset was used to build the classification model, while the testing dataset was used to evaluate model performance.

Sentiment categories were defined as binary classes. A tweet was categorized as positive if it expressed support, agreement, or acceptance of the fuel price increase policy. A tweet was categorized as negative if it reflected rejection, criticism, dissatisfaction, or protest toward the policy.

Manual labeling was performed on the training dataset. Each tweet in the training subset was read and assigned a sentiment label according to the predefined criteria. The remaining subset, referred to as the testing dataset, was not manually labeled. Instead, sentiment labels for

the testing data were generated automatically using the trained classification model.

All labeling and dataset management procedures were performed within RapidMiner Version 10.1 to ensure consistency between preprocessing and modeling environments.

2.5. Feature Extraction

Prior to classification, textual data were transformed into numerical representations. Tokenization was applied to split each tweet into individual word units. This process converted raw text into a structured set of tokens representing candidate features.

Feature weighting was then performed using the Term Frequency–Inverse Document Frequency method. Term Frequency measured the number of times a term appeared in a given tweet, while Inverse Document Frequency quantified the uniqueness of the term across the entire corpus. The TF–IDF weight of term [16]:

$$\text{TF-IDF}(t, d) = \text{TF}(t, d) \times \log\left(\frac{N}{\text{DF}(t)}\right) \quad (1)$$

where $\text{TF}(t, d)$ denotes the frequency of term t in document d , $\text{DF}(t)$ represents the number of documents containing term t , and N is the total number of documents in the corpus. This transformation produced a weighted feature matrix that served as input to the classification algorithm.

2.6. Sentiment Classification

Sentiment classification was performed using the Naïve Bayes algorithm as implemented in RapidMiner Studio version 10.1. The classification procedure was conducted after textual data were transformed into numerical feature vectors through the TF–IDF weighting scheme using the *Process Documents from Data* operator.

Since TF–IDF produces continuous numerical attributes, the Naïve Bayes operator in RapidMiner estimates class-conditional likelihoods under the assumption of a Gaussian distribution for each feature. Therefore, the implemented model corresponds to a Gaussian Naïve Bayes classifier rather than the multinomial variant commonly used for raw term counts. Let $C \in \{c_1, c_2\}$ denote the sentiment classes, where c_1 represents positive sentiment and c_2 represents negative sentiment. Let $X = \{x_1, x_2, \dots, x_p\}$ represent the TF–IDF feature vector of a tweet, where p denotes the total number of extracted terms.

The posterior probability of class C_k given feature vector X is computed according to Bayes' theorem:

$$P(C_k | X) = \frac{P(X | C_k) P(C_k)}{P(X)}. \quad (2)$$

Under the Naïve Bayes conditional independence assumption, the joint likelihood is factorized as:

$$P(X | C_k) = \prod_{i=1}^p P(x_i | C_k). \quad (3)$$

For numerical attributes, each feature likelihood is modeled using the Gaussian density function:

$$P(x_i | C_k) = \frac{1}{\sqrt{2\pi\sigma_{ik}^2}} \exp\left(-\frac{(x_i - \mu_{ik})^2}{2\sigma_{ik}^2}\right), \quad (4)$$

where μ_{ik} and σ_{ik}^2 denote the mean and variance of feature i within class C_k , estimated from the training data.

The prior probability of each class is calculated as:

$$P(C_k) = \frac{N_k}{N}, \quad (5)$$

where N_k is the number of training observations in class C_k and N is the total number of training observations.

For classification, each tweet is assigned to the class that maximizes the posterior probability:

$$\hat{C} = \arg \max_{c_k} \left[\log P(C_k) + \sum_{i=1}^p \log P(x_i | C_k) \right]. \quad (6)$$

A logarithmic transformation is applied to avoid numerical underflow during computation [17]. Model training was conducted using the manually labeled training dataset. The trained model was subsequently applied to the testing dataset using the *Apply Model* operator. All parameter settings were left at RapidMiner default values unless otherwise specified.

2.7. Model Evaluation

Model performance was evaluated using a confusion matrix computed on the testing dataset. The confusion matrix compared predicted sentiment labels with the true labels in the testing data. Let:

TP = True Positive, TN = True Negative, FP = False Positive, FN = False Negative.

$$\begin{aligned} \text{Accuracy} &= \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}}, \\ \text{Precision} &= \frac{\text{TP}}{\text{TP} + \text{FP}}, \\ \text{Recall} &= \frac{\text{TP}}{\text{TP} + \text{FN}}, \\ \text{F1} &= 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}. \end{aligned}$$

All evaluation metrics were generated within RapidMiner Version 10.1 using the classification performance operator. No interpretative analysis of metric values is presented in this section, as performance outcomes are reported separately in the Results section [18].

3. Results and Discussion

This section presents the empirical findings of the study and interprets them in light of the proposed research questions and theoretical framework. The results are organized according to the sequential stages of analysis, beginning with data acquisition and preprocessing, followed by classification outcomes and performance evaluation. Each subsection reports the statistical output and provides an analytical interpretation to contextualize the findings within energy policy sentiment research.

3.1. Crawling Data

This research began with identifying keywords related to #BBM (fuel) using Google Trends. From this activity, five keywords were identified as closely related to #BBM: #hargaBBM, #BLT, #BLTfuel, #BBMnaik, and #cek (#fuel prices, #BLT, #BLTfuel, #fuelprices increases, and #checks). The trend of these keywords is evident in Fig. 2.

Among the identified terms, two hashtags were considered most directly relevant to the fuel price increase policy: #hargabbm and #bbmnaik. These hashtags explicitly refer to fuel prices and price increases, thereby closely aligning with the research objective of capturing public reactions to the policy announcement. Consequently, they were selected as the primary query terms for Twitter data collection.




Fig. 2: Fuel Keywords Trends in 2022

Twitter data crawling was conducted using Google Colab as the computational environment. The retrieval process was restricted to 3–4 September 2022, corresponding to the immediate period following the official fuel price announcement. The detailed distribution of collected tweets is presented in Table 1.

Table 1: Twitter Data Crawling Results

Date	Fuel Prices	Fuel Rises
September 3, 2022	522	618
September 4, 2022	612	251
Sum	1134	869
Total	2003	

Based on Table 1, a total of 2,003 tweets were initially collected and analyzed as raw data. These tweets were retrieved using keyword-based crawling with the hashtags #hargabbm and #bbmnaik. Data collection was conducted over a two-day period, from 3 to 4 September 2022, corresponding to the immediate aftermath of the official policy announcement. This narrow time window was deliberately selected to capture spontaneous and emotionally driven public reactions before discourse became influenced by formal narratives or prolonged media framing.

3.2. Preprocessing Data

Data preprocessing was conducted to enhance the quality and reliability of the Twitter dataset prior to sentiment classification. Several cleaning procedures were systematically applied, including the removal of URLs, hyperlinks, user mentions, emojis, emoticons, special characters, and excessive whitespace. In addition, duplicate tweets containing identical textual content were identified and eliminated to prevent redundancy and bias in the analysis. The data-cleaning process is shown in Fig. 3 and overview of subprocess step is shown in Fig. 4.




Fig. 3: RapidMiner Canvas View of the Data Cleaning Process




Fig. 4: Workflow of Subprocess step

The results of the data cleaning process are presented in Table 2, which displays a side-by-side comparison of tweet content before and after preprocessing. This comparison illustrates how noise elements were eliminated and how the textual data were standardized prior to feature extraction and sentiment classification.

Table 2: The State of Data Before and After Data Cleaning

Operator	Purpose	Command	Amount of Data
Replace (Replace URL 1)	Remove links that are in the middle or end of a sentence Before: Heheh.. This is more on target, so the right people will get it. Stuttgart.. Peralite Prices Rise, Ojol Brothers and Fishermen Get https://t.co/bidD9oRk9r Fuel Social Assistance. #BantuanBBMUntukRakyat #BLTBBMTepatSasaran #HematCermatBBM After: Heheh.. This is more on target, so the right people will get it. Stuttgart.. Peralite Prices Rise, Ojol Brothers and Fishermen Get #Fuel Assistance for the People # BLT Fuel Right on Target #Save Fuel Carefully Fuel Social Assistance	<code>https.* ?space</code>	2003
Replace (Replace URL 2)	Remove links at the beginning of a sentence Before: https://t.co/jPqCrhtRWZ The Government is determined to increase fuel prices when world oil prices are plummeting. #FuelPricesRisePeopleScream After: The Government is determined to increase fuel prices when world oil prices are plummeting. #FuelPricesRisePeopleScream	<code>https.*</code>	2003
Replace (Replace Mention 1)	Remove the mention in the middle or end of a sentence Before: @fajarnugros Astagirullah Na'uzubillahiminzalik #BBMNaik_RakyatMenjerit After: Astagirullah Na'uzubillahiminzalik #FuelPricesRisePeopleScream	<code>@.*?space</code>	2003
Replace (Replace Mention 2)	Removing mentions at the beginning of sentences Before: Workers Provide Three Solutions Related to the Increase in Fuel Prices #workers #fuelpricesincrease #pentaliteincrease @pertamina After: Workers Provide Three Solutions Related to the Increase in Fuel Prices #workers #fuelpricesincrease #pentaliteincrease	<code>@.*</code>	2003
Replace (Replace Simbol)	Removing non-letter and number symbols/characters Before: Somehow the weather is natural this morning, for sure I always pray for you to be enthusiastic in the morning even though I hear that fuel prices are soaring... #fuelpricesincrease #rp10000 #pentalite #morning #morning #morningspirit #galangrambuanarki After: Somehow the natural weather this morning is sure that I always pray for you to be enthusiastic in the morning even though I hear that the price of fuel soars to Rp10000 pentalite morning morningspirit galangrambuanarki	<code>['~?.,". " ##*@()-]</code>	2003
Replace (Replace Emoji)	Removing emojis/emoticons Before: Join the long queue yesterday from 1 o'clock and guess what is really long, it turns out that there is a new fare change from 1430 🚗 o'clock, not the old price eh 😊 bbmincrease After: Join the long queue yesterday from 1 o'clock and guess what is really long, it turns out that there is a new fare change from 1430 o'clock, not the old price eh bbmincrease	<code>([^\x 00-\x 7F]+ @S +)</code>	2003
Trim	Remove excess spaces		2003
Remove Duplicates	Removing the same sentence/tweet		1867

As a result of these preprocessing steps, the dataset was reduced from 2,003 raw tweets to 1,867 unique tweets, indicating that approximately 6.8% of the data consisted of noise or repeated content. This reduction improved the analytical focus by retaining only tweets that conveyed substantive opinions related to the fuel price policy.

The cleaned tweets exhibited distinct linguistic characteristics typical of social media discourse, including informal language, abbreviations, slang, and frequent use of hashtags and mixed Indonesian terms. These features reflect spontaneous public expression and underscore the importance of rigorous preprocessing to ensure accurate sentiment classification.

3.3. Labelling Data

The distribution of sentiment labels in this study reflects public responses to Indonesia's 2022 fuel price policy as captured through Twitter discourse. From the cleaned dataset, 489 tweets were manually labeled and used as training data for the supervised classification model. These tweets were categorized into positive and negative sentiment classes based on explicit expressions of support or opposition toward the fuel price increase. The labeling process ensured that both sentiment categories were represented, allowing the model to learn distinguishing linguistic patterns associated with policy acceptance and rejection.

Examples of tweets labeled as positive or negative are presented in Table 3. The table presents representative samples from each sentiment category, according to the predefined labeling criteria. These examples demonstrate how tweets expressing support or acceptance of the fuel price increase policy were categorized as positive, while tweets containing criticism, rejection, or dissatisfaction were categorized as negative. Presenting these samples provides transparency regarding the operational definition of each sentiment class prior to model training.

Table 3: Example of Data Labeling

Tweets	Label
BUZERRR NKRI HAS BEEN SPREAD which is paid from the results of the people's happiness over suffering	Negative
Fuel Subsidy Diversion: For the Sake of Targeted Subsidies and Anticipation of World Oil Turmoil, Fuel Subsidy Increases Latest Finance news	Positive

Fig. 5 presents the distribution of manually labeled tweets used as training data in the sentiment classification model. A total of 489 tweets were annotated, comprising 69 positive instances (14.1%) and 420 negative instances (85.9%). The chart displays both absolute frequencies and proportional percentages to provide a comprehensive representation of class composition. The visualization highlights the class imbalance within the training dataset, with negative sentiment substantially dominating the labeled corpus. This distribution serves as the empirical foundation for supervised model training, where parameter estimation of the Naïve Bayes classifier is derived exclusively from these manually annotated observations.




Fig. 5: Distribution of Manually Labeled Training Data

3.4. Feature Extraction

Feature extraction was conducted in RapidMiner Studio version 10.1 using a structured processing workflow shown in Fig. 6. The procedure began with importing the dataset using the *Read Excel* operator. The tweet dataset, stored in Excel format, was loaded into the RapidMiner environment to enable further analytical processing. Once imported, the data were available as structured examples for subsequent transformation.




Fig. 6: RapidMiner Canvas View of Feature Extraction Process

After data import, the *Filter Examples* operator was applied to retain only relevant observations for analysis. This filtering step removed records that did not meet the predefined criteria, including empty entries, duplicate records, or tweets outside the designated sentiment categories. The objective of this stage was to ensure data consistency prior to text transformation and feature generation.

The *Nominal to Text* operator was then used to convert attributes with nominal data types into text format. This transformation was necessary because the subsequent text processing operator, *Process Documents from Data*, requires textual input for tokenization and feature extraction. Converting the attribute type ensured compatibility within the RapidMiner processing pipeline. This stage is shown in Fig. 7.




Fig. 7: Overview of Process Documents from Data

The main feature extraction process was executed using the *Process Documents from Data* operator. Within this operator, several sequential text processing steps were performed. Tokenization was applied to split each tweet into individual word units. Case normalization was conducted to standardize all characters into a uniform format. Token filtering was then applied to remove non-informative elements. After these preprocessing steps, word weighting was performed using the Term Frequency–Inverse Document Frequency method. This procedure transformed each tweet into a numerical feature vector, resulting in a document–term matrix that represented the weighted importance of each term across the corpus.

The resulting TF–IDF feature matrix was then used as input for the Naïve Bayes classification model. The Naïve Bayes operator was applied to train the sentiment classification model using the numerically represented training data. The algorithm estimated class probabilities based on the distribution of weighted terms within each sentiment category.

After model training, the *Store Model* operator was used to save the trained Naïve Bayes model. This step enabled the model to be reused for testing, validation, or subsequent implementation without requiring retraining. In addition to saving the model, the processed training dataset was stored using the *Store Data* operator. This preserved the feature-transformed dataset for documentation, reproducibility, and future experimental comparison.

3.5. Classification

A total of 1,378 unlabeled data points were used for testing and classified using the Naïve Bayes Classifier, trained on 489 manually labeled data points. The classification process was also conducted using RapidMiner Studio version 10.1 as the primary analytical platform. The workflow was structured sequentially to ensure consistency between preprocessing outputs and classification inputs. The process is shown in Fig. 8.




Fig. 8: RapidMiner Canvas View at the Classification Stage

The classification began with the *Retrieve Data Latih* operator, which was used to access the previously stored training dataset. This step ensured that the structure of the feature space and term weighting scheme used during model training were aligned with those of the testing dataset.

To unify the feature structure between the training and testing datasets, the *Union* operator was applied. This operator combined the processed testing data with the feature structure derived from the training data, ensuring that both datasets shared the same dimensional feature space. Aligning the feature space was necessary for accurate model application. Following the union process, a second filtering stage was conducted using the *Filter Examples* operator. This step separated the testing data from the training data within the merged structure and removed any unnecessary records prior to model application.

The *Replace Missing Values* operator was applied to handle missing entries arising from feature alignment by imputing a predefined value (e.g., zero), thereby ensuring stable computation during model application. The trained Naïve Bayes model was then loaded using *Retrieve Model* and applied to the testing dataset to produce sentiment predictions; the resulting classification distribution is summarized in Fig. 9.




Fig. 9: Sentiment Classification Result on Testing Data

Fig. 9 presents the distribution of sentiment predictions generated by the Naïve Bayes classifier on the testing dataset. A total of 1,378 tweets were classified, comprising 173 tweets predicted as positive (12.6%) and 1,205 tweets predicted as negative (87.4%). The chart displays both absolute frequencies and proportional percentages to provide a comprehensive representation of model output. The dominance of the negative class is visually evident, reflecting the overall distribution pattern observed during classification. This graphical representation summarizes the predictive distribution produced by the trained model when applied to previously unseen testing data.

3.6. Model Evaluation

The performance of the Naïve Bayes classifier was evaluated using a confusion matrix derived from the training dataset. The confusion matrix summarizes the number of correctly and incorrectly classified tweets across the two sentiment categories, namely positive and negative. True Positive (TP) represents positive tweets that were correctly classified as positive, while True Negative (TN) indicates negative tweets correctly identified as negative. False Positive (FP) refers to negative tweets misclassified as positive, and False Negative (FN) denotes positive tweets incorrectly classified as negative. This matrix provides a comprehensive overview of the classifier's predictive behavior and error distribution.

Table 4: Confusion Matrix of Training Data

	True Negative	True Positive	Class Precision
Pred Negative	395	0	100.00%
Pred Positive	25	69	73.40%
Class Recall	94.05%	100.00%	

Table 4 presents the confusion matrix and classification performance of the Naïve Bayes model evaluated on 489 manually labeled tweets. The model correctly identified 395 negative tweets and 69 positive tweets. A total of 25 negative tweets were incorrectly classified as positive, while no positive tweets were misclassified as negative. The overall classification accuracy achieved was 94.89%. The precision value reached 100.00% for the negative class and 73.40% for the positive class. The recall values were 94.05% for negative tweets and 100.00% for positive tweets. These results summarize the model's predictive performance on the manually annotated evaluation dataset.

Based on Table 4, accuracy and recall can be calculated as follows:

$$\text{Accuracy} = \frac{TP + TN}{TP + FP + FN + TN} = \frac{69 + 395}{69 + 25 + 0 + 395} = 94.89\%$$

$$\text{Precision} = \frac{TP}{TP + FP} = \frac{69}{69 + 25} = 73.40\%$$

$$\text{Recall} = \frac{TP}{TP + FN} = \frac{69}{69 + 0} = 100\%$$

$$\text{F1-score} = \frac{2 \times (\text{Recall} \times \text{Precision})}{\text{Recall} + \text{Precision}} = \frac{2 \times (100 \times 73.40)}{100 + 73.40} = 84.66\%$$

3.7. Discussion

This study analyzed Twitter sentiment toward Indonesia's 2022 fuel price increase and evaluated the performance of a Naïve Bayes classifier for policy-related sentiment classification. Two principal findings emerged. First, public discourse during the immediate post-announcement period was dominated by negative sentiment. Second, the Naïve Bayes model demonstrated strong predictive performance, indicating that probabilistic text classification can reliably identify sentiment polarity in energy policy discourse.

The first research question concerned the dominant sentiment following the policy announcement. The empirical distributions of both the labeled training data and the predicted test data show a clear predominance of negative sentiment. Thus, the hypothesis that negative sentiment would dominate is supported. Substantively, this indicates that the fuel price increase was primarily interpreted as an economic burden rather than a necessary fiscal adjustment.

The second research question examined the effectiveness of the Naïve Bayes classifier. The model achieved high overall accuracy with strong recall and precision values across classes. Therefore, the hypothesis predicting strong classifier performance is also supported. Importantly,

this result suggests that even under informal and noisy linguistic conditions typical of social media, probabilistic generative models remain statistically robust.

The findings align with prior research showing that energy-related policies often generate negative public sentiment, particularly when subsidies affect household purchasing power [19], [20]. Similar to previous sentiment energy studies, this research confirms that digital discourse reflects economically sensitive reactions. Methodologically, the strong performance of Naïve Bayes is consistent with earlier comparative studies showing that simple probabilistic models remain competitive in text classification tasks [16], [21]. While some literature reports superior results from Support Vector Machines or deep neural networks, such advantages often emerge in large-scale corpora [22], [23]. In contrast, the present study uses a short temporal window and domain-specific hashtags, conditions under which simpler generative models may generalize more efficiently [24].

Differences from deep learning studies may also be explained by model–data fit. Complex models require larger balanced datasets to avoid variance inflation. Given the moderate training size and observable class imbalance, Naïve Bayes likely achieved a favorable bias–variance tradeoff.

Several limitations must be acknowledged. First, the temporal scope was restricted to a two-day window. While statistically advantageous for capturing immediate reactions, this design limits inference about the persistence or evolution of sentiment. Second, hashtag-based sampling may introduce selection bias by excluding tweets without specified keywords. The dataset may therefore represent highly engaged users rather than the broader population. Third, the training data exhibited substantial class imbalance. Although performance metrics were strong, imbalance may inflate accuracy and affect minority-class precision stability. Fourth, evaluation relied on a single train–test split rather than k -fold cross-validation. This limits the robustness assessment of model generalizability. Fifth, sentiment was operationalized as a binary variable. This reduces conceptual granularity and may obscure neutral or ambivalent policy evaluations.

4. Conclusion

The study analysed public sentiment towards Indonesia’s increase in fuel prices in 2022 through Twitter data and tested the effectiveness of a Naïve Bayes classifier for sentiment classification on public policy. The main goals were to extract the prevailing sentiment in digital public discourse and also evaluate the statistical robustness of a supervised machine learning method for energy policy analysis.

The results lead to a number of important conclusions. First, overwhelmingly negative sentiment ruled the day on Twitter during the immediate aftermath of the policy announcement, in line with our hypothesis that public reaction would lean toward displeasure. Secondly, the Naïve Bayes classifier achieved good predictive performance which suggests that probabilistic models are still favourable for high-dimensional textual data. Third, TF–IDF representation and Gaussian likelihood estimation resulted in stable discrimination of classes in the presence of informal language and class imbalance. Fourth, this study highlights that short-term data from social media can be an immediate measurable proxy of the public response to economically sensitive policy decisions.

The theoretical contribution of the study lies in its role in tighter integration between applied statistical modeling and public policy analysis and in showing the ongoing relevance of generative probabilistic classifiers for contemporary text analytics. On a more practical basis, the findings suggest that real-time sentiment monitoring can be used by policymakers as an evidence-informed tool to gauge public reaction and adapt communication strategies. Statistical practitioners may mirror this workflow for alternate domains with a need for rapid and interpretable sentiment computation.

Future research should extend the temporal scope, apply cross-validation and comparative modeling techniques, and explore multi-class or longitudinal sentiment dynamics to enhance robustness and generalizability.

CRedit Authorship Contribution Statement

Alia Lestari: conceptualization; methodology; model implementation; visualization; manuscript drafting. **Muhammad Hajarul Aswad:** conceptual refinement; data preprocessing; validation; manuscript review. **Subekti Masri:** validation; manuscript review and editing.

Declaration of Generative AI and AI-Assisted Technologies

Grammarly was used for language refinement and style consistency, and ChatGPT (version 5.2) was used sparingly to enhance sentence clarity and structural coherence during manuscript preparation. These tools were not used to generate data, conduct analysis, or derive scientific conclusions. All academic content, interpretations, and methodological decisions remain the sole responsibility of the authors, and all AI-assisted outputs were reviewed to ensure accuracy and scholarly integrity.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this article.

Funding and acknowledgments

This research was funded by the Litapdimas Research Program of the Institute for Research and Community Service (LP2M) of the State Islamic Institute (IAIN) Palopo, under the 2023 fiscal year research grant scheme. The authors gratefully acknowledge financial support from IAIN Palopo and extend their appreciation to their respective academic institutions for facilitating the completion of this work.

Data and Code Availability

The dataset used in this study consists of publicly available tweets retrieved from the Twitter platform using predefined keywords. Due to Twitter's data redistribution policy, the full text of tweets cannot be publicly shared; however, tweet IDs or metadata can be provided upon reasonable request.

References

- [1] K. Kpodar and B. Liu, "The distributional implications of the impact of fuel price increases on inflation," *Energy Economics*, vol. 108, p. 105909, Apr. 2022. DOI: [10.1016/j.eneco.2022.105909](https://doi.org/10.1016/j.eneco.2022.105909).
- [2] A. Drabo, "Social unrests and fuel prices: The role of macroeconomic, social and institutional factors," International Monetary Fund, IMF Working Paper 2023/228, Oct. 2023, p. 1. DOI: [10.5089/9798400257384.001](https://doi.org/10.5089/9798400257384.001).
- [3] F. Jawadi, D. Bourghelle, P. Rozin, A. I. Cheffou, and G. S. Uddin, "Sentiment and energy price volatility: A nonlinear high frequency analysis," *Energy Economics*, vol. 133, p. 107465, May 2024. DOI: [10.1016/j.eneco.2024.107465](https://doi.org/10.1016/j.eneco.2024.107465).
- [4] S. Saha, "Asymmetric impact of oil price changes on stock prices: Evidence from country and sectoral level data," *Journal of Economics and Finance*, vol. 46, no. 2, pp. 237–282, Apr. 2022. DOI: [10.1007/s12197-021-09559-3](https://doi.org/10.1007/s12197-021-09559-3).
- [5] C. Couharde and E. M. Mouhoud, "Fossil fuel subsidies, income inequality, and poverty: Evidence from developing countries," *Journal of Economic Surveys*, vol. 34, no. 5, pp. 981–1006, Dec. 2020. DOI: [10.1111/joes.12384](https://doi.org/10.1111/joes.12384).

- [6] I. Stupak, M. Mansoor, and C. T. Smith, “Conceptual framework for increasing legitimacy and trust of sustainability governance,” *Energy, Sustainability and Society*, vol. 11, no. 1, p. 5, Mar. 2021. DOI: [10.1186/s13705-021-00280-x](https://doi.org/10.1186/s13705-021-00280-x).
- [7] M. R. Jazuli, K. Roll, and Y. Mulugetta, “A discourse analysis of fuel subsidy reduction: Revisiting the political economy of indonesia’s experiences 1998–2019,” *Humanities and Social Sciences Communications*, vol. 13, no. 1, p. 47, Dec. 2025. DOI: [10.1057/s41599-025-06342-1](https://doi.org/10.1057/s41599-025-06342-1).
- [8] D. Abbonato, *Public sentiments on the fourth industrial revolution: An unsolicited public opinion poll from twitter*, arXiv preprint, Accessed: 2026-02-25, Nov. 2024. <https://arxiv.org/abs/2411.14230v2>.
- [9] X. Ning, R. Li, D. Lan, C. Chen, and Y. Li, “Public sentiment dynamics in policy transitions: A sentiment analysis based on weibo data,” *Frontiers in Public Health*, vol. 13, p. 1692166, Oct. 2025. DOI: [10.3389/fpubh.2025.1692166](https://doi.org/10.3389/fpubh.2025.1692166).
- [10] G. Wang, Y. Liu, and S. Tu, “Discursive use of stability in new york times’ coverage of china: A sentiment analysis approach,” *Humanities and Social Sciences Communications*, vol. 10, no. 1, p. 666, Oct. 2023. DOI: [10.1057/s41599-023-02165-0](https://doi.org/10.1057/s41599-023-02165-0).
- [11] B. N. Nguyen, “A sentiment-based approach to predict energy price volatility using distilroberta and garch models,” *Energy Economics*, vol. 149, Sep. 2025. DOI: [10.1016/j.eneco.2025.108646](https://doi.org/10.1016/j.eneco.2025.108646).
- [12] P. Deniz and T. Stengos, “Electricity prices through the lens of sentiments for the un and the imf: An asymmetric approach for türkiye,” *Borsa Istanbul Review*, vol. 25, no. 4, pp. 756–769, Jul. 2025. DOI: [10.1016/j.bir.2025.03.008](https://doi.org/10.1016/j.bir.2025.03.008).
- [13] A. Priadana and A. A. Rizal, “Sentiment analysis on government performance in tourism during the covid-19 pandemic period with lexicon based,” *CAUCHY: Jurnal Matematika Murni dan Aplikasi*, vol. 7, no. 1, pp. 28–39, Nov. 2021. DOI: [10.18860/ca.v7i1.12488](https://doi.org/10.18860/ca.v7i1.12488).
- [14] B. Agarwal and N. Mittal, “Machine learning approaches for sentiment analysis,” in *Data Mining and Analysis in the Engineering Field*, May 2014, pp. 193–208. DOI: [10.4018/978-1-4666-6086-1.ch011](https://doi.org/10.4018/978-1-4666-6086-1.ch011).
- [15] N. V. Patel and H. Chhinkaniwala, “Investigating machine learning techniques for user sentiment analysis,” *International Journal of Decision Support System Technology*, vol. 11, no. 3, pp. 1–12, 2019. DOI: [10.4018/ijdsst.2019070101](https://doi.org/10.4018/ijdsst.2019070101).
- [16] M. Reusens et al., “Evaluating text classification: A benchmark study,” *Expert Systems with Applications*, vol. 254, p. 124302, Nov. 2024. DOI: [10.1016/j.eswa.2024.124302](https://doi.org/10.1016/j.eswa.2024.124302).
- [17] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. Cambridge, MA: The MIT Press, 2012, 1067 pp.
- [18] K. M. Ting, “Confusion matrix,” in *Encyclopedia of Machine Learning*, 2011, p. 209. DOI: [10.1007/978-0-387-30164-8_157](https://doi.org/10.1007/978-0-387-30164-8_157).
- [19] S. J. Zawadzki, L. Vrieling, and E. van der Werff, “What influences public acceptability of sustainable energy policies? the crucial role of funding and who benefits,” *Energy Research & Social Science*, vol. 87, p. 102468, May 2022. DOI: [10.1016/j.erss.2021.102468](https://doi.org/10.1016/j.erss.2021.102468).
- [20] I. Stadelmann-Steffen and C. Eder, “Public opinion in policy contexts. a comparative analysis of domestic energy policies and individual policy preferences in europe,” *International Political Science Review*, vol. 42, no. 1, pp. 78–94, Jan. 2021. DOI: [10.1177/0192512120913047.10.1177/0192512120913047/suppl_file/supplementary_material_for_public_opinion_in_policy_context.pdf](https://doi.org/10.1177/0192512120913047.10.1177/0192512120913047/suppl_file/supplementary_material_for_public_opinion_in_policy_context.pdf).
- [21] “Effective comparative evaluation of sentiment analysis using paired t-test: A performance study of supervised methods,” *Journal of Southwest Jiaotong University*, vol. 58, no. 5, 2023. DOI: [10.35741/issn.0258-2724.58.5.28](https://doi.org/10.35741/issn.0258-2724.58.5.28).

- [22] D. C. Gkikas and P. K. Theodoridis, “Predicting online shopping behavior: Using machine learning and google analytics to classify user engagement,” *Applied Sciences*, vol. 14, no. 23, Dec. 2024. DOI: [10.3390/app142311403](https://doi.org/10.3390/app142311403).
- [23] Y. C. Lin, S. A. Chen, J. J. Liu, and C. J. Lin, “Linear classifier: An often-forgotten baseline for text classification,” in *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, vol. 2, Jun. 2023, pp. 1876–1888. DOI: [10.18653/v1/2023.acl-short.160](https://doi.org/10.18653/v1/2023.acl-short.160).
- [24] C. Chen et al., “Cross-platform violence detection on social media: A dataset and analysis,” in *Proceedings of the ACM (Conference Proceedings)*, May 2025, pp. 494–498. DOI: [10.1145/3717867.3717877](https://doi.org/10.1145/3717867.3717877).