



Topic-Driven Sentiment Analysis of Indihome Customer Opinions using Indobert and Latent Dirichlet Allocation

Yusza Reditya Murti^{1*}, Aryo Nitiyoga², Herry Irawan³, Aisyah Fadillah Al-Adawiyah⁴, Edward Kandia Tenggehi⁵

^{1,3,4}Digital Business Study Program, School of Economics and Business, Telkom University, Bandung, Indonesia

^{2,5}Management Study Program, School of Economics and Business, Telkom University, Bandung, Indonesia

Correspondence: E-mail: yuszaa@telkomuniversity.ac.id

ABSTRACT

This study analyses customer satisfaction with IndiHome, Indonesia's dominant fixed broadband provider, using a Topic-Driven Sentiment Analysis (TDSA) framework applied to Twitter data. A corpus of 5,674 tweets was collected via the Twitter API between June and July 2021 and labelled into three sentiment classes. IndoBERT was fine-tuned on the labelled corpus, achieving a testing accuracy of 80% and validation accuracy of 90%. Latent Dirichlet Allocation (LDA) was then applied independently to each sentiment subcorpus, producing nine topic clusters subsequently mapped onto Zeithaml et al.'s five customer satisfaction dimensions. Results show that negative sentiment dominates at 53%, driven by three distinct concerns: network connectivity failures, unresponsive complaint handling, and accumulated customer frustration. Positive sentiment centres on network reliability and fair service treatment. The product and service features dimension emerges as the primary satisfaction driver across all sentiment classes, while attribution of service failure surfaces in neutral rather than negative topics, indicating that dissatisfied customers express frustration emotionally before seeking resolution, and attribution of service success appears in positive sentiment alongside perceptions of equity and fairness. The TDSA framework demonstrates that per-class topic modelling produces more actionable service insights than document-level sentiment classification alone, offering a replicable analytical approach for ISP customer feedback monitoring in Indonesia.

ARTICLE INFO

Article History:

Submitted/Received

11 October 2025

First Revised 20 October 2025

Accepted 03 November 2025

First Available online

07 November 2026

Publication Date

01 December 2026

Keyword:

sentiment analysis; IndoBERT;
latent Dirichlet allocation;
customer satisfaction; internet
service provider; Indonesia

1. INTRODUCTION

The rapid expansion of internet access in Indonesia has reshaped how people work, study, and consume services. Based on data from the Indonesian Internet Service Providers Association (APJII), the number of internet users in Indonesia reached 196.7 million in the 2019–2020 period, representing an increase of approximately 8.9% from the previous year (Wearesocial, 2021). This growth, accelerated significantly by the shift to remote work and online learning during the COVID-19 pandemic, has placed fixed broadband providers under intense scrutiny. Demand surged, but so did expectations and complaints.

IndiHome, operated by PT Telkom Indonesia, occupies a dominant position in the Indonesian fixed broadband market. According to the World Bank, IndiHome commands approximately 87% of the country's fixed broadband internet market, making it by far the most widely used provider in the country (Annur, 2021). Despite its scale, IndiHome has consistently attracted a high volume of public criticism on social media. By the end of 2020, the company reached 8.2 million subscribers, falling short of its own target of 8.5 million users (Pebrianto, 2021). At the same time, a survey by PwC found that customer satisfaction with internet networks in Indonesia declined during 2020, a trend that is plainly visible in the volume and tone of user-generated content on Twitter (Pebrianto, 2021).

Social media platforms, particularly Twitter, have become important channels through which customers express frustration, approval, and suggestions about the services they use. Kotler and Keller (2016) note that user reviews serve as meaningful indicators of perceived product performance and customer experience. Given that IndiHome attracts a large and vocal user base on Twitter, the platform provides a rich source of unsolicited customer feedback which captures real-time service perceptions without the framing effects of formal surveys. Sentiment analysis, defined as the computational identification of opinions and evaluative judgements in text, offers a structured way to extract insight from this otherwise unstructured data (Ding et al., 2018).

Prior studies on IndiHome-related sentiment have established a consistent finding: negative sentiment dominates. Hernandi and Wibowo (2021) employed a BiLSTM model on IndiHome Twitter data and focused specifically on detecting negative sentiment. Hermansyah and Sarno (2021) applied aspect-based sentiment analysis on IndiHome product and service reviews to assess quality dimensions. More recently, Qolbu et al. (2025) compared the performance of Naïve Bayes, SVM, and IndoBERT on IndiHome Twitter data, demonstrating that IndoBERT outperformed the other models with an accuracy of 0.96. More recently still, Siregar et al. (2025) applied IndoBERT to geo-tagged IndiHome tweets on platform X, extending the analysis to geographic sentiment distribution. These studies have made meaningful contributions, but they share a common limitation: they classify sentiment at the document level without integrating topic modelling or mapping findings to a customer satisfaction framework. Sentiment polarity alone, without topic-level attribution, offers limited actionable value to service managers.

This study addresses that gap by combining IndoBERT-based sentiment classification with per-class Latent Dirichlet Allocation (LDA) topic modelling, an approach characterised here as Topic-Driven Sentiment Analysis (TDSA) and mapping the resulting topic distributions onto Zeithaml et al.'s five dimensions of customer satisfaction (Wilson et al., 2016). The findings are expected to offer IndiHome and similar Indonesian ISPs a more granular picture of customer dissatisfaction, grounded in what customers discuss on social media rather than what survey instruments are designed to capture.

2. LITERATURE REVIEW

2.1. Customer Satisfaction and Social Media Feedback

Customer satisfaction is broadly understood as the outcome of comparing a customer's perceived experience of a product or service to their prior expectations. When performance falls short of expectations, dissatisfaction follows; when it meets or exceeds them, satisfaction is the result (Kotler & Keller, 2016). Zeithaml et al. identify five factors that shape how customers arrive at these evaluations: the features of the product or service itself, the emotions experienced during consumption, attributions of success or failure in service delivery, perceptions of equity and fairness in how they were treated, and the influence of other consumers or social referents (Wilson et al., 2016). These factors are particularly relevant in the context of subscription-based services like fixed broadband internet, where customers regularly interact with the service and have clear expectations about reliability, speed, and responsiveness.

What has changed substantially in recent years is where dissatisfaction gets expressed. Social media platforms, Twitter in particular, have become the dominant public channel for customer feedback, replacing or supplementing formal complaint mechanisms. Unlike survey responses, which are structured and solicited, social media posts reflect spontaneous and unfiltered opinions. Sethna et al. (2017) observe that user-generated content on social platforms carries significant influence not only as a record of individual experience, but as a collective signal of brand perception. For service providers operating at the scale of IndiHome, monitoring this stream of public opinion is no longer optional, it is a practical necessity for understanding where service quality falls below customer expectations.

2.2. Text Mining and Sentiment Analysis

Text mining refers to the process of extracting meaningful patterns and knowledge from unstructured or semi-structured textual data (Ding et al., 2018). It sits at the intersection of data mining and natural language processing (NLP), drawing on computational methods to make large volumes of text amenable to systematic analysis. NLP itself is the subfield of artificial intelligence concerned with enabling machines to understand, interpret, and generate human language, a task that is considerably more complex than it might appear, given the ambiguities, contextual dependencies, and informal variations that characterise everyday written language, particularly on social media (Ghosh et al., 2012).

Sentiment analysis is one of the most widely applied branches of text mining. It refers to the computational identification of subjective information in text such as opinions, evaluations, attitudes, and emotional stances, typically categorised as positive, negative, or neutral (Ding et al., 2018; Birjali et al., 2021). In the context of customer feedback, sentiment analysis allows researchers and practitioners to process volumes of data that would be impractical to read manually. A corpus of tens of thousands of tweets, for instance, can be classified and summarised in a fraction of the time required for human annotation. The practical value of this is well-established across domains ranging from product reviews to public health monitoring, and its application to internet service provider feedback is a natural extension of this broader trend.

It is worth noting, however, that document-level sentiment classification, assigning a single polarity label to each piece of text, captures only part of the picture. Knowing that 53% of

IndiHome-related tweets carry negative sentiment tells us something important, but it does not tell us whether customers are frustrated about internet speed, billing, customer service responsiveness, or something else entirely. This distinction matters for any practical intervention. Sentiment without topic attribution has limited diagnostic value.

2.3. Deep Learning for Indonesian Text: From BERT to IndoBERT

Machine learning approaches to text classification generally involve training a model on labelled examples so that it can generalise to new, unseen text. Early methods relied on feature engineering, constructing explicit representations of text using techniques such as TF-IDF weighting or bag-of-words models, before passing these features to classifiers like Naïve Bayes or Support Vector Machines. These methods work reasonably well for many classification tasks, but they treat words as independent units and struggle to capture the contextual relationships that give language its meaning. The word "slow" means something different in "slow but reliable connection" compared to "slow, unreliable, and overpriced," yet classical models have difficulty making this distinction.

Neural networks and deep learning changed this. By stacking multiple layers of computational units, deep learning models learn hierarchical representations of input data, from low-level patterns at early layers to abstract, task-relevant features at deeper ones (Jurafsky & Martin, 2019). Recurrent architectures such as LSTM introduced the ability to model sequential dependencies in text, and convolutional models contributed spatial pattern recognition. Both represented genuine advances over classical methods, though they still required substantial labelled data and were sensitive to preprocessing decisions.

The introduction of transformer-based models marked a more fundamental shift. BERT (Bidirectional Encoder Representations from Transformers) is a pre-trained language model developed by Google that learns deep bidirectional contextual representations of text by training on massive corpora using masked language modelling and next sentence prediction (Devlin et al., 2018). Unlike earlier models that processed text left-to-right or right-to-left, BERT reads in both directions simultaneously, giving it a richer understanding of context. Crucially, BERT can be fine-tuned for specific downstream tasks, including sentiment classification, with relatively small amounts of task-specific labelled data, which makes it highly practical for applied research.

IndoBERT extends this architecture to the Indonesian language. Developed as part of the IndoNLU benchmark project, IndoBERT follows the BERT-base configuration with 12 hidden layers, 768-dimensional representations, 12 attention heads, and feed-forward hidden layers of 3,072 dimensions (Wilie et al., 2020). It was pre-trained on over 220 million Indonesian words drawn from Wikipedia, national news outlets including Kompas, Tempo, and Liputan6, and the Indonesian Web Corpus. This domain-specific pre-training gives IndoBERT a substantial advantage over multilingual models when processing Indonesian text, particularly informal Twitter language that includes slang, abbreviations, and code-switching, all of which are common in IndiHome-related tweets.

2.4. Topic Modelling and the TDSA Framework

Topic modelling is an unsupervised machine learning technique designed to discover the latent thematic structure within a collection of documents. Rather than assigning predefined categories to text, topic models infer recurring patterns of co-occurring words from the data itself, treating each document as a mixture of topics and each topic as a distribution over

words (Hagen, 2018). Latent Dirichlet Allocation (LDA), proposed by Blei et al. (2003), remains the most widely used method in this family. LDA has been applied across a broad range of domains, from political discourse and healthcare to e-commerce and service quality evaluation, and its ability to extract interpretable themes from large corpora without requiring labelled training data makes it well-suited to social media analysis.

When applied to the full corpus of a dataset, LDA reveals what topics are discussed overall. This is informative, but in the context of sentiment-tagged data, a more targeted application is possible and considerably more revealing. By running LDA separately on each sentiment class: negative, positive, and neutral, the resulting topics become anchored to a specific emotional valence. Negative-sentiment topics show what customers are complaining about; positive-sentiment topics show what they appreciate. This approach, which can be characterised as Topic-Driven Sentiment Analysis (TDSA), integrates the outputs of sentiment classification and topic modelling into a unified analytical pipeline rather than treating them as two separate exercises (Ozyurt & Akcayol, 2021).

The distinction between standard sentiment analysis and TDSA is not merely technical, it is interpretive. Standard sentiment analysis answers the question of how customers feel. TDSA answers both how they feel and what specific aspects of the service are the source of those feelings. For service providers, this is the difference between knowing that customers are dissatisfied and knowing that their dissatisfaction is concentrated around network reliability or billing disputes. When the topic outputs from TDSA are further mapped onto a structured framework of customer satisfaction dimensions, such as the five-factor model proposed by Zeithaml et al. (Wilson et al., 2016), the computational findings become directly translatable into management-relevant conclusions. This is the analytical approach adopted in the present study.

3. MATERIAL AND METHODS

3.1. Research Framework

This study employs a Topic-Driven Sentiment Analysis (TDSA) framework that integrates IndoBERT-based sentiment classification with per-class LDA topic modelling to examine IndiHome customer feedback on Twitter. As illustrated in Figure 1, the pipeline proceeds sequentially from data collection and preprocessing through manual labelling, IndoBERT sentiment classification, and LDA topic modelling applied independently to each sentiment subcorpus, concluding with a mapping of the resulting topic distributions onto Zeithaml et al.'s five customer satisfaction dimensions (Wilson et al., 2016). Running LDA within each sentiment class, rather than on the full corpus, is the defining feature of TDSA: it anchors each topic to a specific emotional valence and enables direct interpretation of what customers discuss when they feel a particular way about IndiHome.




Figure 1 Topic-Driven Sentiment Analysis Framework

3.2. Data Collection and Corpus Construction

Indonesian-language tweets containing the keyword "IndiHome" were collected between 12 June and 28 July 2021 via the Twitter API using Python in a Jupyter Notebook environment, yielding an initial corpus of 10,272 tweets. The corpus was constructed through keyword-based data retrieval; the unit of analysis is the tweet as a textual document, and the relevant population is defined as the set of publicly available Indonesian-language tweets matching the keyword and time window criteria, a framing consistent with corpus-based NLP research (Ding et al., 2018). Following removal of advertisements, retweets, and spam, the working corpus comprised 5,674 tweets used for labelling, model training, and analysis.

3.3. Preprocessing and Dataset Labelling

Preprocessing was applied in two layers corresponding to the two downstream tasks. For IndoBERT fine-tuning, only surface cleaning was performed: case folding, removal of URLs, usernames, hashtags, numbers, symbols, and punctuation, normalisation of informal

abbreviations using a custom dictionary, and tokenisation. Stopword removal and stemming were deliberately excluded at this stage, as IndoBERT uses subword tokenisation (WordPiece) and computes contextual embeddings across the full input sequence, aggressive preprocessing disrupts this contextual signal and has been shown empirically to reduce classification performance in Indonesian transformer models (Khairani et al., 2024; Devlin et al., 2018). For LDA topic modelling, applied after sentiment classification to each sentiment subcorpus separately, stopwords removal and Sastrawi-based stemming were additionally applied, as LDA is a bag-of-words model where stopwords introduce noise and stemming improves topic coherence. Examples of each preprocessing step are presented in Tables 1 and 2. Following preprocessing, the 5,674 tweets were manually labelled into three sentiment categories: negative (0), neutral (2), and positive (1), with representative examples shown in Table 3.

Table 1. Surface cleaning preprocessing examples (applied to all pipeline stages)

Step	Original tweet	Result
Case folding	Jangan KUMAT MULU deh INTERNET INDIHOME !!!	jangan kumat mulu deh internet indihome !!!
Data cleaning	@IndiHome Aduh kepikiran ya :) internetnya mati mulu	aduh kepikiran ya internetnya mati mulu
Tokenisation	ya ampun ini gimana sih. Internet down, katanya perbaikan jaringan tapi. Terus gw wfh gimana ini kalo 3 hari wifi ga nyala.	'ya', 'ampun', 'ini', 'gimana', 'sih', 'internet', 'down', 'katanya', 'perbaikan', 'jaringan', 'terus', 'gw', 'wfh', 'gimana', 'ini', 'kalo', 'hari', 'wifi', 'ga', 'nyala'

Table 2. LDA-specific preprocessing examples (applied to sentiment subcorpora only, not prior to IndoBERT)

Step	Original tweet	Result
Normalisation	capek bgt ama indihome	'capek', 'banget', 'sama', 'indihome'
Stopword removal	paling enak emang pakai indihome	'enak', 'pakai', 'indihome'
Stemming	bagi yg kepikiran mau make, inget ya siapin ekstra 150k buat kuota jangan sekali2 harepin provider ini bakal lancar tiap saat	'bagi', 'yang', 'pikir', 'mau', 'pakai', 'ingat', 'siapin', 'ekstra', 'kuota', 'jangan', 'harepin', 'provider', 'bakal', 'lancar', 'tiap', 'saat'

Table 3. Dataset labelling examples

Tweet	Sentiment label
Indihome gua jelek banget lemot wifi nya dari kemaren, kapan diperbaiki??	Negative (0)
Saran buat admin indihome tolong bales keluhan pelanggan dong, udah lapor dari tadi ga ada respon	Neutral (2)
Cinta banget sama indihome, internet aman lancar buat WFH sebulan ini	Positive (1)

3.4. IndoBERT Sentiment Classification

IndoBERT, pre-trained on over 220 million Indonesian words and configured with 12 hidden layers, 768-dimensional representations, and 12 attention heads (Wilie et al., 2020), was fine-tuned on the labelled corpus for three-class sentiment classification. Fine-tuning involved adding a classification layer to the pre-trained encoder and training the full network on the task-specific data at a low learning rate, preserving the general linguistic knowledge acquired during pre-training while adapting the model to the IndiHome domain (Devlin et al., 2018). The corpus was split 70:15:15 for training, validation, and testing. Three epoch settings were evaluated at batch size 32 consistent with standard practice for fine-tuning BERT-scale models (Osinga, 2018) and learning rate 2e-5; epoch 5 was selected as the final configuration based on the best balance between training accuracy (96%) and validation accuracy (90%), with epoch 10 showing signs of overfitting (99% training, 89% validation). The research framework is visualised in Figure 3.




Figure 2. Pre-Training & Fine Tuning

3.5. LDA Topic Modelling and Satisfaction Mapping

Following sentiment classification, the labelled corpus was partitioned into negative, neutral, and positive subcorpora. LDA was fitted independently to each subcorpus using Python in Google Colab, with three topics per subcorpus, producing nine topic clusters in total. Topic outputs were visualised as word clouds (Figures 8–10) and subsequently mapped onto Zeithaml et al.'s five customer satisfaction dimensions (Wilson et al., 2016) through close reading of topic word distributions and representative tweets, following interpretive conventions of qualitative content analysis applied in conjunction with computational methods (Blei et al., 2003).

3.6. Evaluation Metrics

Model performance is reported using four metrics derived from the confusion matrix:

$$Accuracy = (TP + TN) / (TP + FP + FN + TN) \quad (1)$$

$$Precision = TP / (TP + FP) \quad (2)$$

$$Recall = TP / (TP + FN) \quad (3)$$

$$F1 - score = 2 \times (Precision \times Recall) / (Precision + Recall) \quad (4)$$

Validation accuracy (90%) refers to performance on the validation split during training; testing accuracy (80%) refers to performance on the held-out test set and is the primary generalisation measure. The confusion matrix is presented in Figure 6.




Figure 6. Confusion Matrix Diagram

4. RESULTS AND DISCUSSION

4.1. Sentiment Classification Results

The preprocessing stage reduced the raw corpus from 10,272 tweets to 5,674 usable data points after the removal of advertisements, retweets, and spam content (Figure 4). The cleaned corpus was then manually labelled and split for IndoBERT fine-tuning. Two split ratios were tested, 60:20:20 and 70:15:15, and the results are summarised in Table 4. The 70:15:15 ratio produced a higher validation accuracy (90% versus 83%) and was adopted as the final configuration for all subsequent training runs.

Table 4. Dataset split ratio comparison

Ratio	Training accuracy	Validation accuracy
60:20:20	96%	83%
70:15:15	96%	90%

Three epoch settings were evaluated using a batch size of 32 and a learning rate of $2e-5$, as detailed in Table 5. Epoch 5 yielded the most balanced outcome, training accuracy of 96% and validation accuracy of 90%, with a smaller gap between the two compared to epoch 10 (99% training, 89% validation), which showed signs of overfitting. Epoch 4, while close in validation accuracy (88%), did not reach the same level of generalisation. On this basis, epoch 5 was selected as the final model configuration.

Table 5. IndoBERT epoch comparison results

Epoch	Training accuracy	Validation accuracy
4	98%	88%
5	96%	90%
10	99%	89%

The final model was evaluated on the held-out test set, producing an overall testing accuracy of 80% as reported in the classification report (Figure 5). Performance varied across the three sentiment classes. Negative sentiment classification was strongest, with precision 89%, recall 84%, and F1-score 86%. Positive sentiment performed reasonably well at precision 84%, recall 80%, and F1-score 82%. Neutral sentiment proved the most difficult to classify accurately, returning precision 56%, recall 70%, and F1-score 62%. The lower performance on neutral tweets is consistent with the inherent ambiguity of this class, tweets in this category often lack strong polarity markers and may contain a mix of complaint and inquiry without a clear emotional stance. The confusion matrix (Figure 6) confirms this pattern: the model correctly classified 356 negative, 229 positive, and 96 neutral instances, with the largest proportion of misclassifications concentrated in the neutral class.

The overall sentiment distribution on the test corpus is presented in Figure 7. Negative sentiment accounts for 53% of tweets, positive for 29%, and neutral for 18%. This distribution indicates that most of the IndiHome-related discourse on Twitter during the study period was unfavourable, with more than half of all tweets carrying negative sentiment. The dominance of negative sentiment is consistent with findings from prior studies on IndiHome Twitter data (Hernandi & Wibowo, 2021; Qolbu et al., 2025) and aligns with the documented decline in Indonesian internet service customer satisfaction during this period (Pebrianto, 2021).

4.2. Topic Modelling Results

LDA was applied independently to each of the three sentiment subcorpora, producing three topics per class, nine topics in total across the full corpus. The word cloud visualisations for negative, positive, and neutral sentiment are presented in Figures 8, 9, and 10 respectively. Table 6 consolidates all nine topics into a single TDSA summary, mapping each topic to its dominant keywords and the corresponding Zeithaml et al. customer satisfaction dimension.

Table 6. TDSA summary: topic distributions mapped to customer satisfaction dimensions

Sentiment	Topic	Dominant keywords	Zeithaml dimension
Negative	Topic 1	jelek, sinyal, bikin, capek, emosi, error, kesel	Consumer emotions
Negative	Topic 2	internet, mati, koneksi, wifi, ganggu, bayar, pagi	Product & service features

Negative	Topic 3	lemot, parah, jaringan, wifi, internet, rumah, kemarin	Product & service features
Neutral	Topic 1	internet, mati, wifi, rumah, cek, bantu, kemarin	Attribution of service failure
Neutral	Topic 2	wifi, anak, rumah, coba, please, kenapa, suka	Product & service features
Neutral	Topic 3	bayar, ganggu, lemot, sinyal, bikin, jam, nonton	Product & service features
Positive	Topic 1	lancar, internet, belajar, rumah, paket, gampang, betah	Product & service features
Positive	Topic 2	aman, cepat, baik, kasih, terima, pasang, anak	Attribution of service success + Perceptions of equity & fairness
Positive	Topic 3	bagus, kenceng, seneng, terbaik, keren, wifi, mantap	Consumer emotions

Note: Keywords represent the highest-probability terms per topic extracted via LDA (n_components=3, random_state=42) applied independently to each sentiment subcorpus (negative n=2,998; positive n=1,661; neutral n=1,014). The keyword "indihome" was excluded from all topic outputs as it appears uniformly across all documents. Dimension mapping follows Zeithaml et al. as operationalised in [Wilson et al. \(2016\)](#).

Three patterns emerge from Table 6 that are worth noting before the discussion that follows. First, the product and service features dimension is the most pervasive, appearing across all three sentiment classes and accounting for six of the nine topic clusters, it manifests as connectivity disruptions and slow speeds in the negative and neutral classes, and as reliable performance and ease of use in the positive class. Second, consumer emotions appear at both ends of the sentiment spectrum: negative Topic 1 captures accumulated frustration and emotional distress (jelek, emosi, capek, kesel), while positive Topic 3 reflects satisfaction and pride in service quality (bagus, seneng, terbaik, mantap), an asymmetry where negative emotional expression is more intense and more varied than its positive counterpart. Third, the attribution dimension is present in both the neutral and positive classes but takes opposite forms: neutral Topic 1 reflects customers seeking service recovery with no resolution in sight (cek, bantu, mati), while positive Topic 2 reflects customers crediting IndiHome for responsive and reliable delivery (cepat, baik, terima, aman). Particularly, attribution of service failure does not emerge as a dominant negative topic in its own right, suggesting that when IndiHome fails, customers express frustration emotionally (Topic 1) or descriptively around the service failure itself (Topics 2 and 3), rather than directing blame explicitly at the provider.

4.3. Discussion: Customer Satisfaction Dimensions

The TDSA results presented in Table 6 reveal that IndiHome customer satisfaction is shaped primarily by three of Zeithaml et al.'s five dimensions: product and service features, consumer emotions, and attribution of service success or failure. The product and service features dimension is the most pervasive, appearing across all three sentiment classes, connectivity failures and slow speeds dominate negative topics, while speed, stability, and reliable performance for work-from-home and streaming use cases drive positive expression. This confirms network performance as the central determinant of how customers evaluate IndiHome, consistent with ISP service quality studies that use Twitter-based NLP methods ([Rintyarna et al., 2022](#)).

The consumer emotions dimension is concentrated at both ends of the sentiment spectrum but takes strikingly different forms. In the negative class, Topic 1 captures intense, accumulated frustration, terms such as *jelek*, *emosi*, *capek*, and *kesel* reflect customers who have reached a breaking point rather than reporting a single incident. In the positive class, Topic 3 reflects pride and satisfaction (*bagus*, *seneng*, *terbaik*, *mantap*), though the emotional vocabulary here is less varied and less forceful than its negative counterpart, suggesting that positive experiences generate milder emotional expression than negative ones. The attribution dimension presents a more nuanced picture than prior studies have suggested. Attribution of service failure does not emerge as a dominant negative topic; instead, when IndiHome fails, customers express their dissatisfaction either emotionally (Negative Topic 1) or by describing the service failure directly (Negative Topics 2 and 3), rather than directing explicit blame at the provider. Attribution of service failure does, however, surface in neutral Topic 1, where customers seek resolution, reporting outages, requesting technician visits, and asking for confirmation, but without the emotional intensity of the negative class. Conversely, attribution of service success appears in positive Topic 2, where customers acknowledge responsive and reliable delivery (*cepat*, *baik*, *terima*, *aman*), often alongside perceptions of equity and fairness. The social influence dimension did not emerge from any of the LDA outputs, consistent with the first-person, experience-focused character of the Twitter corpus.

These findings suggest that IndiHome's most pressing service improvement priorities lie in network reliability and complaint responsiveness, two dimensions where the gap between customer expectation and perceived performance is widest. The TDSA framework makes this actionable in a way that document-level sentiment classification alone cannot: it identifies not just that customers are dissatisfied, but specifically what they are dissatisfied about and through which emotional and relational channels that dissatisfaction is expressed.

5.CONCLUSION

This study applied a Topic-Driven Sentiment Analysis (TDSA) framework to 5,674 IndiHome-related tweets, combining IndoBERT sentiment classification with per-class LDA topic modelling to analyse customer satisfaction across positive, negative, and neutral sentiment classes. The IndoBERT model achieved a testing accuracy of 80% and validation accuracy of 90%, with negative sentiment dominating the corpus at 53%, followed by positive at 29% and neutral at 18%. Topic modelling revealed that negative sentiment encompasses three distinct concerns: network connectivity failures, accumulated customer frustration, and slow or disrupted service, while positive sentiment centres on network reliability, responsive service delivery, and fair treatment. Product and service features emerged as the most pervasive satisfaction driver across all sentiment classes, with attribution of service failure surfacing in neutral rather than negative topics, reflecting customers who seek resolution before emotional expression escalates. The TDSA approach demonstrates that integrating per-class topic modelling with sentiment classification produces more actionable insights than document-level sentiment analysis alone, offering IndiHome a diagnostic tool to identify not just the volume of dissatisfaction but its specific service dimensions. Future research should consider expanding the dataset across multiple platforms, balancing sentiment class distribution, and evaluating alternative topic modelling methods such as BERTopic to further improve topic coherence on short-form social media text.

6. REFERENCES

- Annur, C. M. (2021). *IndiHome mendominasi pasar internet fixed broadband di Indonesia*. Katadata.
- Birjali, M., Kasri, M., & Beni-Hssane, A. (2021). A comprehensive survey on sentiment analysis: Approaches, challenges and trends. *Knowledge-Based Systems*, 226, 107134.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022.
- Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. *ArXiv Preprint ArXiv:1810.04805*.
- Ding, Z., Li, Z., & Fan, C. (2018). Building energy savings: Analysis of research trends based on text mining. *Automation in Construction*, 96, 398–410.
- Ghosh, S., Roy, S., & Bandyopadhyay, S. K. (2012). A tutorial review on text mining algorithms. *International Journal of Advanced Research in Computer and Communication Engineering*, 1(4), 120–126.
- Hagen, L. (2018). Content analysis of e-petitions with topic modeling: How to train and evaluate LDA models? *Information Processing & Management*, 54(6), 1292–1307.
- Hernandi, M. K., & Wibowo, S. A. (2021). Sentiment analysis implementation for detecting negative sentiment towards IndiHome in Twitter using bidirectional long short-term memory. *2021 IEEE International Conference on Communication, Networks and Satellite (COMNETSAT)*, 310–315.
- Hermansyah, R., & Sarno, R. (2021). Sentiment analysis based on quality aspects in effort to improve quality of IndiHome product and services PT Telkom Indonesia Tbk. *3rd International Conference on Computer Science and Computational Intelligence (ICCSCI 2021)*. Atlantis Press.
- Jurafsky, D., & Martin, J. H. (2019). *Speech and language processing* (3rd ed.). Prentice Hall.
- Khairani, U., Mutiawani, V., & Ahmadian, H. (2024). Pengaruh tahapan preprocessing terhadap model IndoBERT dan IndoBERTweet untuk mendeteksi emosi pada komentar akun berita Instagram. *Jurnal Teknologi Informasi dan Ilmu Komputer*, 11(4), 887–894.
- Kotler, P., & Keller, K. L. (2016). *A framework for marketing management* (6th ed.). Pearson.
- Osinga, D. (2018). *Deep learning cookbook: Practical recipes to get started quickly*. O'Reilly Media.
- Ozyurt, B., & Akcayol, M. A. (2021). A new topic modeling based approach for aspect extraction in aspect based sentiment analysis: SS-LDA. *Expert Systems with Applications*, 168, 114231.
- Pebrianto, F. (2021). *Survei PwC: Kepuasan pelanggan internet Indonesia menurun selama pandemi*. Tempo.
- Qolbu, A. A., Fitriyati, N., & Inayah, N. (2025). Performa Naïve Bayes, SVM, dan IndoBERT pada analisis sentimen Twitter IndiHome dengan strategi penanganan data tidak seimbang. *Jurnal Fourier*, 14(1), 29–44.

- Rintyarna, B. S., Kuswanto, H., Sarno, R., Rachmaningsih, E. K., Rachman, F. H., Suharso, W., & Cahyanto, T. A. (2022). Modelling service quality of internet service providers during COVID-19: The customer perspective based on Twitter dataset. *Informatics*, 9(1), 11.
- Sethna, B. N., Hazari, S., & Bergiel, B. J. (2017). Influence of user generated content in online shopping: Impact of gender on purchase behaviour, trust, and intention to purchase. *International Journal of Electronic Marketing and Retailing*, 8(4), 344–371.
- Siregar, R. A., Fitriyani, F., Komariah, K., Lesmana, T., & Ruldeviyani, Y. (2025). Sentiment analysis of IndiHome service based on geo location using the BERT model on platform X. *Jurnal Teknik Informatika (JUTIF)*, 6(1), 233–242.
- Wearesocial. (2021). *Digital 2021: The latest insights into the state of digital*. We Are Social.
- Wilie, B., Vincentio, K., Winata, G. I., Cahyawijaya, S., Li, X., Lim, Z. Y., & Purwarianti, A. (2020). IndoNLU: Benchmark and resources for evaluating Indonesian natural language understanding. *ArXiv Preprint ArXiv:2009.05387*.
- Wilson, A., Zeithaml, V., Bitner, M. J., & Gremler, D. (2016). *Services marketing: Integrating customer focus across the firm* (3rd European ed.). McGraw-Hill.