Comparative Analysis of Transformer-Based Models for Zero-Shot Abstractive Text Summarization

Ishita Chaturvedi[1], Pooja Gupta[1,2], Rajiv Singh[1,2]*

[1]Department of Computer Science, Banasthali Vidyapith, Tonk, Rajasthan 304022, India
[2]Centre for Artificial Intelligence, Banasthali Vidyapith, Tonk, Rajasthan 304022, India
{ishitachaturvedi5,poojagupta2291, jkrajivsingh*}@gmail.com

Abstract. Abstractive text summarization has become an important task in Natural Language Processing (NLP) due to the rapid growth of textual data across various domains. Although several transformer-based models have been developed for text summarization, their performance has not been compared systematically in a unified zero-shot setting. This study presents a comparative evaluation of three pretrained transformer-based models, DistilBART, FLAN-T5, and ProphetNet on CNN/DailyMail dataset in a zero-shot setting. The models generate summaries without any task-specific fine-tuning, and their performance is evaluated using ROUGE metrics. In addition, TF-IDF Cosine Similarity, Sentence-BERT Similarity, and Jaccard Similarity are used to compare the lexical, semantic, and token-level similarity of the generated summaries. The results show that DistilBART achieves the best overall ROUGE performance by maintaining a good balance between precision, recall, content coverage and structural coherence. FLAN-T5 generates concise summaries with moderate accuracy, whereas ProphetNet achieves higher recall but low precision because it includes additional information. The similarity analysis shows that ProphetNet preserves semantic meaning more effectively, while DistilBART maintains a balanced lexical and structural representation. Overall, the results indicate that model architecture and pertaining strategy affect summarization quality and can help in selecting suitable models for practical text summarization tasks.

Keywords: Abstractive text summarization, zero-shot learning, transformer-models.

1 Introduction

In today's digital era, volume of textual data is increasing rapidly due to the continuous generation of content through news articles, research papers and online platforms. Due to this growth, manual extraction of relevant and concise information has become difficult and time-consuming. To solve this problem, automatic text summarization becomes an important area for Natural Language Processing (NLP), its purpose is to convert large documents into short and meaningful summaries without losing any essential information [1].

Automatic summarization is widely used in information retrieval, document analysis and content understanding systems, which gives users important information fast

© The Author(s) 2026
A. Shastri et al. (eds.), *Proceedings of the Fourth Workshop on Applications of Computational Intelligence and Interactive Systems (WACIIS 2026)*, Advances in Intelligent Systems Research 217,
https://doi.org/10.2991/978-94-6239-832-0_9

and efficiently. It removes the unnecessary details of the information and preserves the main idea [1].

Recent developments in deep learning have significantly improved in the field of abstractive text summarization. In particular, Transformer architecture has become highly effective because of their self-attention mechanism which helps the models to capture the contextual relationships and long-range dependencies within text [2]. Pretrained Transformer models have further improved language understanding and generation by learning general linguistic representations from large-scale corpora [3]. Encoder-decoder architectures such as BART and T5 have also been used for various text summarization tasks [4]. DistilBART [5], FLAN-T5 [6], and ProphetNet [7] use different approaches to improve text summarization using techniques like knowledge distillation, instruction tuning, and future n-gram prediction, respectively.

However, many existing research evaluate these models on task-specific fine-tuning. The evaluation does not directly indicate how well pretrained models can perform summarization without any task specific fine tuning. In a zero-shot setting, input is directly provided to a pretrained model, and it generates a summary without being fine-tuned on the summarization dataset [6]. This setting is particularly useful for practical applications where task-specific training may be costly or not available.

To address this issue, this study provides a unified zero-shot comparison of DistilBART, FLAN-T5, and ProphetNet using the CNN/DailyMail dataset. The models are evaluated with ROUGE metrics along with precision, recall, and F1-score. In addition, TF-IDF cosine similarity, Sentence-BERT similarity, and Jaccard similarity are also used to examine lexical, semantic, and token-level relationships among the generated summaries.

The main contributions of this study are as follows:

- A common zero-shot evaluation framework is developed to compare three transformer-based summarization models without task-specific fine-tuning.
- The models are evaluated using ROUGE metrics.
- The generated summaries are compared using lexical, semantic, and token-level similarity measures.
- The study compares precision, recall, content coverage, and semantic consistency across transformer architectures.

2 Related Work

Recent advancements in abstractive text summarization have been significantly influenced by transformer-based architectures, particularly encoder–decoder models. Among these, three important models- DistilBART, FLAN-T5 and ProphetNet use different approaches for text generation. These models are suitable for zero-shot because they use three different approaches for sequence-to-sequence generation: knowledge distillation, instruction tuning, and future n-gram prediction.

DistilBART is the lightweight version of BART model which is made from knowledge distillation technique. Its main goal is to reduce the size of the model without affecting its performance. It follows the encoder-decoder architecture in which the encoder understands the text and the decoder generate the summary. Research on

Transformer-based summarization has shown that knowledge distillation can reduce the computational cost of large pretrained models while maintaining text generation capabilities [5]. Further studies indicate that DistilBART is effective as well as generates relevant and coherent summaries, which make it suitable for practical applications. Therefore, DistilBART shows the balance between model size, cost, and summarization quality.

FLAN-T5 is the improved version of T5 which is based on instruction tuning, which means the model performs the task by understanding the natural language instructions like "summarize the text" [8]. Instruction tuning improves the ability of pretrained language models to perform different tasks based on the instructions [6]. This approach is very useful in zero-shot settings, because models perform tasks without any training. The FLAN-T5 family has shown strong zero-shot performance on different NLP tasks, which makes it suitable for performing tasks without any task-specific fine tuning [8].

ProphetNet uses a different approach which introduces future n-gram prediction. In this approach instead of predicting one word, it predicts multiple future words, which helps in understanding long-term context. The original ProphetNet framework uses future n-gram prediction and n-stream self-attention to help model to get future tokens during sequence generation [7]. Theoretically it improves coherence, but in practical results it has been observed that the model includes extra content. Due to this, recall goes high, but precision remains low, and summaries include unnecessary information.

By comparing these three models, DistilBART focuses on efficiency and balanced performance. FLAN-T5 provides generalization and flexibility. ProphetNet introduces a new prediction mechanism. Thus, the three models provide a useful base for comparison because of their different architectures and pretraining strategies.

Existing research was mainly focused on fine-tuned models. However, the performance of these models on zero-shot settings may differ because of their different pretrained objectives and task-generalization capabilities. In zero-shot setting there is still limited detailed comparison.

Moreover, there is still limited knowledge of how transformer-based models balance coherence with content coverage, and precision with recall, particularly when compared with different architectures. In addition, only ROUGE evaluation may not fully capture the lexical, semantic, and token-level similarities among summaries generated by different models. This highlights the need for a systematic evaluation of these models within a unified zero-shot framework, which is the main focus of this study. Therefore, this study compares DistilBART, FLAN-T5, and ProphetNet under same zero-shot conditions using ROUGE-based evaluation along with TF-IDF cosine similarity, Sentence-BERT similarity, and Jaccard similarity.

3 Proposed Framework

3.1 Data Collection

This study uses a standard benchmark dataset by which fair and comparable evaluation becomes possible. Primary dataset CNN/DailyMail has been used, which is a large-scale summarization corpus. In these datasets human written summaries are also

available with news articles. Table 1 represents real-world news domain and is widely accepted benchmark for abstractive summarization evaluation [9].

Table 1: Dataset description.

Dataset	Domain	Input	Reference Output
CNN/DailyMail	News Articles	Full News Article	Human-written Summary

3.2 Data Preprocessing

Before summarization, the raw documents are preprocessed to provide clean and structured input to the models. This step raw text is converted into clean and model- friendly format. This process includes multiple steps:

- Noise removal: HTML tags, special characters and irrelevant symbols were removed.
- Tokenization: Sentences and words were split into meaningful tokens.
- Input formatting: Articles and summaries were converted into model-specific format (encoder-decoder input structure).

The main purpose of this study is to maintain consistency, and to provide clean and structured input to model which keeps inference stable.
This preprocessing operation can be represented in Equation 1 as

$$\mathrm{X_p} = \mathrm{P(X)} \tag{1}$$

where X represents the raw input document, P represents the preprocessing function and X_p represents the preprocessed input document

3.3 Model Selection and Zero-Shot Setting

Zero-shot setting means model is directly evaluated without any fine-tuning on any specific dataset. In this approach model generates output on the basis of its pretraining and general knowledge. The main objective is to check how much model generalizes without any task-specific training. It is helpful in real world scenario where separate training is not possible for every domain.

In this framework three transformer-based pretrained models namely DistilBART, FLAN-T5, and ProphetNet were selected. These models were evaluated on zero-shot setting, where dataset is not specifically fine-tuned. Models perform inference directly on the basis of their inference, given in Equation 2.

$$\hat{S} = M_\theta(X) \tag{2}$$

where X is Input Document, M_θ is Pretrained Model and $\hat{S}$ denotes Generated Summary
Now, the generated summary is represented as a sequence of output tokens, given in Equation 3

$$\hat{S} = \{ y1, y2, y3, \ldots., yn\} \tag{3}$$

where $\hat{S}$ denotes generated summary, $y_i$is output token generated at the ith decoding step and n is the total number of generated tokens in the summary

3.4 Evaluation Metrics

ROUGE metrics were used for text summarization evaluation. It measures content overlap between generated and reference summaries and mainly follows recall-based approach [10]. ROUGE-1 evaluates word-level similarity, ROUGE-2 measures phase-level similarity, and ROUGE-L evaluates sequence-based similarity [10]. Overall, ROUGE is used to assess the summarization quality on the basis lexical matching

3.5 Similarity Analysis

Besides ROUGE evaluation, similarity analysis is performed on models to examine how closely the generated summaries match their reference summaries in terms of lexical, semantic, and token-level similarity.
Three similarity measures are used-

TF-IDF (Term Frequency Inverse Document Frequency) Cosine Similarity- It measures the lexical similarity between generated and reference summaries [11], given in Equation 4.

$$Cosine\ (A,B) = \frac{A.B}{||\ A||\ ||B||} \tag{4}$$

where, $A\ and\ B$ are the TF-IDF vectors of the reference and generated summaries.

Sentence-BERT Similarity: It measures the semantic similarity between generated and reference summaries [12], given in Equation 5.

$$Similarity\ (A,B) = \frac{E_A\ .\ E_B}{||E_A\ ||\ ||E_B||} \tag{5}$$

where, E_A andE_B are the Sentence-BERT embeddings of the reference and generated summaries.

Jaccard Similarity: It measures the token-level overlap between generated and reference summaries [13], given in Equation 6.

$$J\ (A,B) = \frac{|A \cap B|}{|A\ \cup B|} \tag{6}$$

where, A and B represent the sets of tokens in reference and generated summaries.

4 Results and Discussions

(CNN)The Palestinian Authority officially became the 123rd member of the International Criminal Court on Wednesday, a step that gives the court jurisdiction over alleged crimes in Palestinian territories. The formal accession was marked with a ceremony at The Hague, in the Netherlands, where the court is based. The Palestinians signed the ICC's founding Rome Statute in January, when they also accepted its jurisdiction over alleged crimes committed "in the occupied Palestinian territory, including East Jerusalem, since June 13, 2014." Later that month, the ICC opened a preliminary examination into the situation in Palestinian territories, paving the way for possible war crimes investigations against Israelis. As members of the court, Palestinians may be subject to counter-charges as well. Israel and the United States, neither of which is an ICC member, opposed the Palestinians' efforts to join the body. But Palestinian Foreign Minister Riad al-Malki, speaking at Wednesday's ceremony, said it was a move toward greater justice. "As Palestine formally becomes a State Party to the Rome Statute today, the world is also a step closer to ending a long era of impunity and injustice," he said, according to an ICC news release. "Indeed, today brings us closer to our shared goals of justice and peace." Judge Kuniko Ozaki, a vice president of the ICC, said acceding to the treaty was just the first step for the Palestinians. "As the Rome Statute today enters into force for the State of Palestine, Palestine acquires all the rights as well as responsibilities that come with being a State Party to the Statute. These are substantive commitments, which cannot be taken lightly," she said. Rights group Human Rights Watch welcomed the development. "Governments seeking to penalize Palestine for joining the ICC should immediately end their pressure, and countries that support universal acceptance of the court's treaty should speak out to welcome its membership," said Balkees Jarrah, international justice counsel for the group. "What's objectionable is the attempts to undermine international justice, not Palestine's decision to join a treaty to which over 100 countries around the world are members." In January, when the preliminary ICC examination was opened, Israeli Prime Minister Benjamin Netanyahu described it as an outrage, saying the court was overstepping its boundaries. The United States also said it "strongly" disagreed with the court's decision. "As we have said repeatedly, we do not believe that Palestine is a state and therefore we do not believe that it is eligible to join the ICC," the State Department said in a statement. It urged the warring sides to resolve their differences through direct negotiations. "We will continue to oppose actions against Israel at the ICC as counterproductive to the cause of peace," it said. But the ICC begs to differ with the definition of a state for its purposes and refers to the territories as "Palestine." While a preliminary examination is not a formal investigation, it allows the court to review evidence and determine whether to investigate suspects on both sides. Prosecutor Fatou Bensouda said her office would "conduct its analysis in full independence and impartiality." The war between Israel and Hamas militants in Gaza last summer left more than 2,000 people dead. The inquiry will include alleged war crimes committed since June. The International Criminal Court was set up in 2002 to prosecute genocide, crimes against humanity and war crimes. CNN's Vasco Cotovio, Kareem Khadder and Faith Karimi contributed to this report.

Figure 1: Input text from CNN/DailyMail dataset

The Palestinian Authority becomes the 123rd member of the International Criminal Court . The move gives the court jurisdiction over alleged crimes in Palestinian territories . Israel and the United States opposed the Palestinians' efforts to join the court . Rights group Human Rights Watch welcomed the move, saying it was a step toward justice .

Figure 2: Generated summary from DistilBART

palestinian authority officially became 123rd member international criminal court wednesday step gives court jurisdiction alleged crimes palestinian territories formal accession marked ceremony in the hague netherlands

Figure 3: Generated summary from FLAN-T5

the palestinian authority officially became the 123rd member of the international criminal court on wednesday a step that gives the court jurisdiction over alleged crimes in palestinian territories. the formal accession was marked with a ceremony at the hague in the netherlands where the court is based the palestinians signed the iccs founding rome statute in january when they also accepted its jurisdiction over alleged crimes committed in the occupied palestinian territory including east jerusalem since june 13 2014.

Figure 4: Generated summary from ProphetNet

Table 2: Average ROUGE scores of DistilBART on CNN/DailyMail dataset

Metric	Average Precision	Average Recall	Average F1
ROUGE-1	**0.3147**	**0.4746**	**0.3741**
ROUGE-2	0.1458	0.2257	0.1748
ROUGE-L	0.2449	0.3715	0.2916

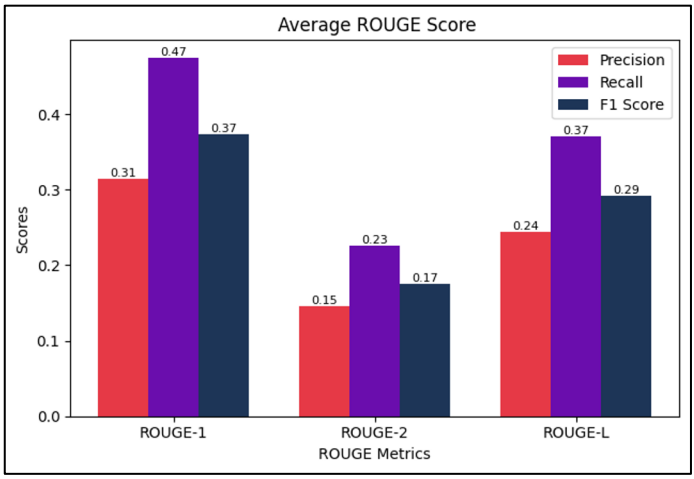

Figure 5: Average ROUGE scores of DistilBART on CNN/DailyMail dataset

Table 3: Average ROUGE scores of FLAN-T5 on CNN/DailyMail dataset

Metric	Average Precision	Average Recall	Average F1
ROUGE-1	**0.2749**	**0.3536**	**0.3053**
ROUGE-2	0.1176	0.1638	0.1352
ROUGE-L	0.2301	0.2891	0.2530

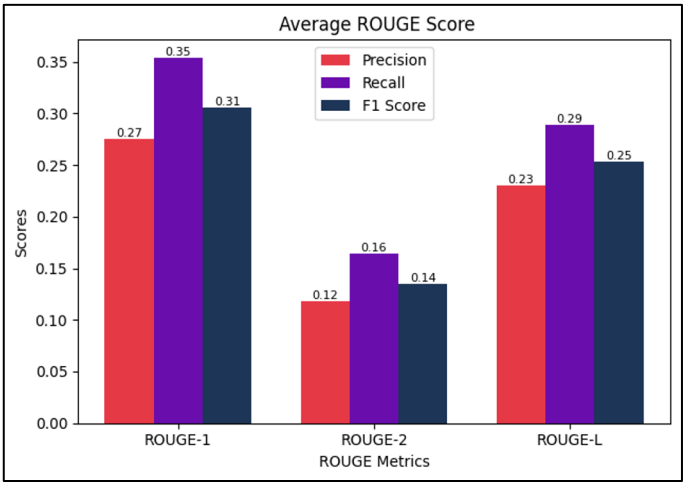

Figure 6: Average ROUGE scores of LAN-T5 on CNN/DailyMail dataset

Table 4: Average ROUGE Scores of ProphetNet on CNN/DailyMail dataset

Metric	Average Precision	Average Recall	Average F1
ROUGE-1	**0.2569**	**0.4693**	**0.3301**
ROUGE-2	0.0686	0.1363	0.0906
ROUGE-L	0.1962	0.3141	0.2186

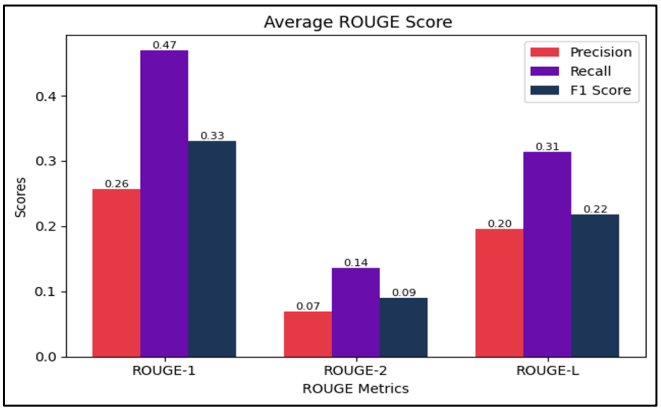

Figure 7: Average ROUGE scores of ProphetNet on CNN/DailyMail dataset

Three summarization models—DistilBART, FLAN-T5, and ProphetNet—have been compared using ROUGE metrics on the CNN/DailyMail dataset shows in Figure 1. These three models' average Rouge scores are defined in Tables 2, 3, and 4, correspondingly. Models were tested on ROUGE-1, ROUGE-2 and ROUGE-L, where precision, recall, and F1 scores were considered and presented in Figure 5, 6 and 7 accordingly. The resulting summaries from the three datasets are displayed in Figures 2, 3, and 4, correspondingly.

It has been observed by that DistilBART gives consistently best results compared to FLAN-T5 and ProphetNet across all ROUGE metrics defined in Table 5. Its F1-scores- ROUGE-1 (0.3741), ROUGE-2 (0.1748), and ROUGE-L (0.2916) indicate that this model effectively maintains content relevance and structural coherence.

Table 5: Comparative analysis of models

Aspect	DistilBART	FLAN-T5	ProphetNet
Overall Performance	High	Moderate	Low
ROUGE-1 (Content)	Balanced	Moderate	High Recall, Low Precision
ROUGE-2 (Coherence)	Strong	Average	Weak
ROUGE-L (Structure)	Strong	Moderate	Weak
Precision-Recall Balance	Well-Balanced	Moderate	Imbalanced
Summary Quality	Best	Acceptable	Lower

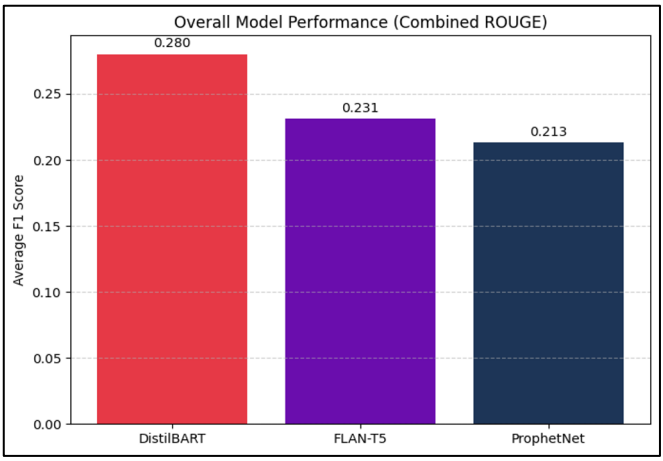

Figure 8: Overall performance comparison based on average ROUGE F1-scores

Table 6: Comparative ROUGE F1 scores of all models (CNN/DailyMail)

Model	ROUGE-1 F1	ROUGE-2 F1	ROUGE-L F1
DistilBART	**0.3741**	**0.1748**	**0.2916**
ProphetNet	0.3301	0.0906	0.2186
FLAN-T5	0.3053	0.1352	0.2530

Comparative results clearly show in Table 6 and Figure 8 that DistilBART performs better than other models. Its knowledge distillation approach makes it fast and efficient which is suitable for practical approach. ProphetNet gives second best performance. Its ROUGE-1 recall is relatively high, which shows that it covers the most content coverage of document. However, low ROUGE-2 is comparatively low, indicates less phrase level consistency and sometimes generates extra content. FLAN-T5 shows moderate performance. Due to instruction tuning it generalizes well on different tasks, but it couldn't maintain precision and context relevance like DistilBART and ProphetNet. Overall ranking:$DistilBART > ProphetNet > FLAN-T5$. The results show in Table 7 that summarization performance depends on the pertaining strategy of models. DistilBART gives the most balanced performance.

Table 7: Qualitative comparison of models on CNN/DailyMail dataset

Model	Architecture	Dataset	Strength
DistilBART	Distilled seq2seq	CNN/DailyMail	Balanced performance
FLAN-T5	Instruction-based	CNN/DailyMail	Flexible
ProphetNet	N-gram prediction	CNN/DailyMail	High Recall

Table 8: TF-IDF cosine similarity

	DistilBART	FLAN-T5	ProphetNet
DistilBART	1.0000	0.3451	0.5744
FLAN-T5	0.3451	1.0000	0.5495
ProphetNet	0.5744	0.5495	1.0000

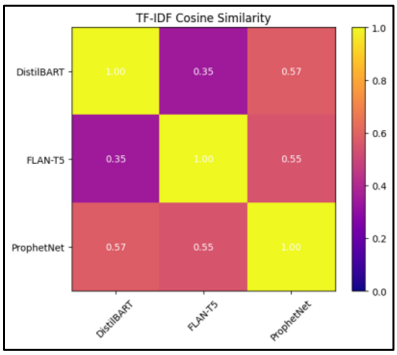

Figure 9: TF-IDF cosine similarity

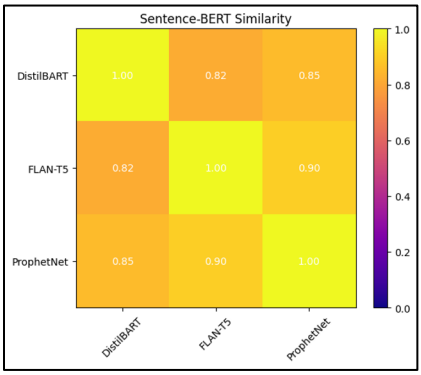

Figure 10: Sentence-BERT similarity

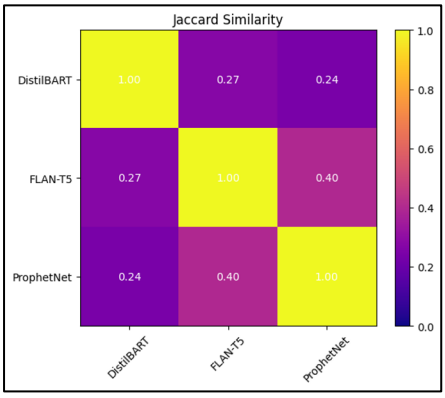

Figure 11: Jaccard similarity

Table 9: Sentence-BERT similarity

	DistilBART	FLAN-T5	ProphetNet
DistilBART	1.0000	0.8163	0.8530
FLAN-T5	0.8163	1.0000	0.9013
ProphetNet	0.8530	0.9013	1.0000

Table 10: Jaccard similarity

	DistilBART	FLAN-T5	ProphetNet
DistilBART	1.0000	0.2745	0.2405
FLAN-T5	0.2745	1.0000	0.3966
ProphetNet	0.2405	0.3966	1.0000

Three metrics were used in similarity analysis to increase the reliability of the comparison: TF-IDF Cosine Similarity, Sentence-BERT Similarity, and Jaccard Similarity, which are described in Tables 8, 9, and 10. Furthermore, these data are displayed with a confusion matrix in Figures 9, 10, and 11, accordingly. These metrics measures lexical, semantic, and token-level overlap respectively defined in Table 11.

Table 11: Comparative Summary of Similarity Metrics

Metric	Similarity Level	Highest Pair (Score)	Observation
TF-IDF Cosine Similarity	Moderate	ProphetNet (0.5477)	Moderate lexical overlap, similar wording patterns
Sentence-BERT Similarity	High	ProphetNet (0.9013)	Strong semantic alignment, meaning almost same
Jaccard Similarity	Low	DistilBART (0.3966)	Low token overlap, abstractive nature reflected

ProphetNet is better choice, when requirement is high and strong context preservation, because its semantic similarity scores are high. These models are particularly useful in news summarization and research content. FLAN-T5 is a good option for generating short and concise summaries as it removes unnecessary details, resulting in compact outputs without losing the core meaning. DistilBART is good choice for balanced performance. It generates readable summaries with moderate lexical similarity. This model does not perform aggressive paraphrasing or generates very long summaries.

5 Conclusion and Future Work

Overall analysis shows that abstractive summarization models use different approaches to process and generate text, but in most cases, semantic consistency is maintained. TF-IDF results show moderate lexical overlap, whereas Sentence-BERT confirms strong semantic alignment across models. The low Jaccard similarity scores clearly highlight that token-based metrics alone are not sufficient for evaluating abstractive summarization. Therefore, it is important to use a combination of multiple metrics for a comprehensive evaluation. A consistent pattern is observed where ProphetNet shows high similarity across all metrics, particularly in Sentence-BERT similarity. This indicates that ProphetNet have strong contextual alignment and generate summaries with highly similar meanings, even if the wordings are differing slightly. DistilBART demonstrates comparatively balanced behavior. In TF-IDF similarity, it suggests that it follows similar lexical patterns. However, its semantic similarity is relatively lower; indicating that the model performs simplification and compression, leading to changes in structure and phrasing. FLAN-T5 follows a clearly different pattern. Its TF-IDF similarity is slightly lower, while Sentence-BERT scores remain high. This indicates that FLAN-T5

performs aggressive paraphrasing; significantly altering words and sentence structure while preserving the overall meaning. Jaccard similarity consistently shows relatively low scores, providing an important insight. It shows that token-level overlaps are minimal, which is expected because abstractive models do not repeat exact words. The future research can explore few-shot learning, hybrid transformer frameworks, and human-based evaluation methods to provide a more comprehensive analysis.

References

1. Nenkova, A., McKeown, K.: A Survey of Text Summarization Techniques. Mining Text Data, Boston, MA, USA: Springer, pp. 43–76 (2012).
2. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention Is All You Need. Advances in Neural Information Processing Systems, vol. 30 (2017).
3. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. in Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 4171–4186 (2019).
4. Lewi, M. et al.: BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension," in Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 7871–7880 (2020).
5. Shleifer, S., Rush, A.M.: Pre-trained Summarization Distillation. arXiv:2010.13002 (2020).
6. Wei, J., et al.: Finetuned Language Models Are Zero-Shot Learners. in International Conference on Learning Representations (ICLR) (2022).
7. Qi, W., Yan, Y., Gong, Y., Liu, D., Duan, N., Chen, J., Zhang, R., Zhou, M.: ProphetNet: Predicting Future N-gram for Sequence-to-Sequence Pre-training. in Findings of the Association for Computational Linguistics: EMNLP 2020, pp. 2401–2410 (2020).
8. Chung, H.W., et al.: Scaling Instruction-Finetuned Language Models. arXiv:2210.11416 (2022).
9. Hermann, K.M., Kocisky, T., Grefenstette, E., Espeholt, L., Kay, W., Suleyman, M., Blunsom, P.: Teaching Machines to Read and Comprehend. in Advances in Neural Information Processing Systems, vol. 28 (2015).
10. Lin, C.-Y.: ROUGE: A Package for Automatic Evaluation of Summaries. in Text Summarization Branches Out, pp. 74–81, Barcelona, Spain (2004).
11. Salton, G., Buckley, C.: Term-weighting approaches in automatic text retrieval. Information Processing & Management, vol. 24, no. 5, pp. 513–523 (1988).
12. Reimers, N., Gurevych, I.: Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. in Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, pp. 3982–3992, Hong Kong, China (2019).
13. Gomaa, M., Fahmy, A.: A Survey of Text Similarity Approaches. International Journal of Computer Applications, vol. 68, no. 13, pp. 13–18 (2013).

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (http://creativecommons.org/licenses/by-nc/4.0/), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

