

# Summarize Smart: Domain-Specific Transformers for Legal and Financial Texts

MODUGU NIKHITHA<sup>1</sup>, PAPINENI RAMASUBBARAO<sup>2</sup>, PATTI TEJASWINI<sup>3</sup>, RAMISETTI PAVAN<sup>4</sup>, MRS.K.PRAVEENA<sup>5</sup>

*Electronics & communication engineering, Chalapathi Institute of Engineering & Technology, LAM, Guntur- Andhra Pradesh<sup>1,2,3,4</sup>.*

<sup>5</sup>Assistant Professor *Electronics & communication engineering, Chalapathi Institute of Engineering & Technology, LAM, Guntur*

**Abstract**— Legal and financial documents are often lengthy, dense, and filled with specialized language, posing a challenge for efficient comprehension and decision-making. Manual summarization is time-consuming, error-prone, and inconsistent, emphasizing the need for automated systems that can deliver concise, accurate, and context-aware summaries. This paper presents a domain-adapted Natural Language Processing (NLP) approach for summarizing legal and financial texts using advanced transformer-based architectures. The proposed framework leverages fine-tuned versions of PEGASUS and T5 models, tailored to understand domain-specific structures and terminology through curated legal-financial corpora. A hybrid summarization pipeline is introduced, combining extractive and abstractive strategies to preserve factual accuracy and maintain logical flow. Additionally, the model integrates named entity recognition (NER) and an attention-guided scoring mechanism to retain critical content such as case references, monetary values, and contractual clauses. The system is benchmarked using evaluation metrics including ROUGE, BLEU, and BERTScore, and compared against conventional models like LexRank, TextRank, and BERTSUM. Results show superior performance in generating coherent, relevant, and domain-accurate summaries. The architecture is scalable, language-independent, and well-suited for applications such as legal analytics, compliance reporting, financial review, and intelligent contract summarization, offering a powerful solution for enhancing document understanding in data-intensive industries.

**Keywords**— Natural Language Processing (NLP), Transformer Models, Legal Summarization, Financial Document Analysis, Hybrid Summarization, PEGASUS, T5, Named Entity Recognition (NER), ROUGE, BLEU, BERTScore.

## I. INTRODUCTION

The exponential growth of legal and financial documentation in the digital era has created a significant bottleneck in knowledge processing and decision-making across sectors such as law, banking, auditing, and compliance. These documents—ranging from contracts and case law to financial disclosures and audit reports—are often long, structurally complex, and rich in domain-specific terminology. Manually reviewing such documents is not only time-consuming but also prone to inconsistency and human error, posing serious risks in high-stakes decision environments. To address these challenges, automated text

summarization has emerged as a promising solution. Summarization systems aim to extract the most relevant information from a document and present it in a concise, readable form. However, generic summarization models frequently fall short when applied to legal and financial texts, as they lack the contextual understanding and terminological precision required to accurately capture critical facts, clauses, and figures.

Recent advances in Natural Language Processing (NLP), particularly transformer-based models such as BERT, T5, BART, and PEGASUS, have significantly improved the capabilities of automatic summarization. These models have shown strong performance in general-domain tasks, yet they often struggle with domain adaptation due to the unique linguistic and structural attributes of legal and financial documents. Summaries generated by such models may contain hallucinated content or omit essential legal and financial entities, making them unsuitable for real-world use without further adaptation. This paper proposes a hybrid summarization framework that integrates both extractive and abstractive techniques, fine-tuned specifically for the legal and financial domains. By leveraging transformer architectures—namely PEGASUS and T5—along with named entity recognition (NER) and attention-guided relevance scoring, our system is designed to retain factual accuracy, preserve domain-specific entities, and ensure the coherence and completeness of summaries. The proposed solution is evaluated using benchmark datasets and industry-standard metrics including ROUGE, BLEU, and BERTScore.

The exponential growth of digital information in legal and financial domains has created a pressing demand for intelligent systems capable of extracting meaningful summaries from complex documents. Legal contracts, court judgments, audit reports, and financial disclosures are often lengthy, verbose, and packed with specialized terminology. Professionals in these domains are burdened with reviewing large volumes of text to identify critical insights, a process that is both time-consuming and prone to human error. Automating the summarization process is, therefore, essential to improve productivity, reduce costs, and enhance decision-making accuracy.

Natural Language Processing (NLP) has emerged as a powerful tool for automating text analysis tasks such as classification, named entity recognition, and summarization. Among these, text summarization has garnered increasing attention, particularly with the advent of transformer-based models such as BERT [1], T5 [2], and PEGASUS [3], which have significantly outperformed traditional recurrent and extractive models. However, these models are typically pre-trained on general-domain corpora and often fail to produce coherent or accurate summaries when applied to domain-specific documents like legal filings or financial statements [4].

In recent years, research efforts have focused on adapting these models to specialized domains using techniques such as domain-specific fine-tuning [5], hybrid summarization approaches [6], and knowledge-aware models that incorporate named entity recognition and clause segmentation [7]. These adaptations have shown promise, but challenges remain in preserving the factual accuracy and context sensitivity required in legal and financial summaries.

In this paper, we propose a domain-adapted hybrid NLP model that leverages both extractive and abstractive summarization techniques, customized for legal and financial text. The model is fine-tuned on curated datasets from publicly available sources and employs attention-based mechanisms to prioritize critical entities and facts. Experimental results show a significant improvement over baseline models in terms of ROUGE, BLEU, and BERTScore metrics. Our approach not only enhances the quality of summaries but also addresses key domain-specific challenges such as terminology consistency, clause linkage, and numerical value retention.

The contributions of this paper are threefold:

- A novel hybrid summarization architecture fine-tuned for legal and financial domains.
- Integration of entity-aware attention mechanisms for improved factual consistency.
- Comprehensive evaluation against standard benchmarks and real-world datasets.

## II. RELATED WORKS

Text summarization has been a long-standing task in Natural Language Processing (NLP), with methods traditionally categorized into extractive and abstractive approaches. Extractive methods select and combine salient sentences from the input text [1], while abstractive methods aim to generate novel sentences that paraphrase the content [2]. While early extractive models like TextRank and LexRank achieved moderate success in general-purpose summarization [3], they often failed to capture the logical and semantic coherence necessary for domain-specific documents.

The advent of transformer-based models marked a significant shift in NLP. BERT [4], though not designed for generative tasks, laid the groundwork for deep contextual understanding. Later models like T5 [5], BART [6], and

PEGASUS [7] demonstrated remarkable performance on benchmark datasets such as CNN/DailyMail and XSum. However, these models were pre-trained on general corpora and struggled when applied to complex legal or financial texts due to domain mismatch.

To address this gap, researchers have developed domain-specific language models, such as LEGAL-BERT [8] for legal documents and FinBERT [9] for financial texts. These models improved domain adaptation through pretraining on relevant corpora but were primarily optimized for classification and named entity recognition tasks rather than summarization. Chalkidis et al. [8] also highlighted the challenge of summarizing legal cases, where preserving legal terminology and hierarchical clause structure is critical.

Recent efforts have explored fine-tuning general-purpose summarization models on domain-specific datasets. Zhong et al. [10] adapted BERT for legal text summarization by integrating clause segmentation and sentence-level relevance ranking. Similarly, Elgandy et al. [11] proposed a hybrid model combining extractive and abstractive techniques to improve factual consistency and fluency in summarizing legal statutes and contracts.

In the financial domain, summarization has been applied to tasks such as earnings report analysis and SEC filings [12], often focusing on extracting financial highlights or sentiment-bearing sentences. However, these models rarely address full document-level summarization and tend to omit contextual dependencies across sections.

While prior work has made progress in adapting NLP models for legal and financial summarization, challenges such as factual accuracy, terminology preservation, and logical coherence remain largely unsolved. Our work builds upon these efforts by proposing a hybrid transformer-based model fine-tuned on carefully curated domain-specific corpora, incorporating attention-based entity preservation and sentence relevance ranking to improve the quality of both extractive and abstractive summaries.

### 2.1. Existing System

Existing summarization systems have shown impressive performance on generic text data such as news articles and Wikipedia entries. However, when applied to legal and financial documents, they often fall short in terms of contextual understanding, domain specificity, and factual correctness.

Most existing models can be broadly categorized into extractive, abstractive, and hybrid summarization systems.

#### 2.1.1 Extractive Summarization Systems

Extractive models identify and rank key sentences from the original text based on syntactic and semantic features. Algorithms like TextRank [1], LexRank [2], and more recently, BERTSum [3] have been employed for extractive summarization. These models tend to preserve grammatical correctness since they copy exact text, but they often lack coherence and can miss out on synthesizing new or implicit information — which is critical in legal and financial contexts where implicit logical connections are common.

#### 2.1.2 Abstractive Summarization Systems

Abstractive methods, such as T5 [4], BART [5], and PEGASUS [6], use encoder-decoder architectures to

generate summaries in natural language. These models are capable of paraphrasing and sentence restructuring, making them suitable for producing fluent and human-like summaries. However, without domain-specific training, these models often generate hallucinated facts or omit critical information, especially when dealing with highly specialized terminology or legal structures [7].

### 2.1.3 Domain-Specific Models

To address these shortcomings, several domain-specific models have been introduced. **LEGAL-BERT** [8] and **FinBERT** [9] are transformer models pre-trained on legal and financial corpora respectively. While effective in classification and information retrieval tasks, they are not specifically optimized for summarization. Attempts have been made to fine-tune generic models on small legal/financial datasets, but the lack of annotated summarization datasets in these domains limits their performance and generalizability [10].

### 2.1.4 Challenges in Existing Systems

Despite these advancements, the following limitations persist in current systems:

- **Lack of factual accuracy** in abstractive outputs.
- **Inability to preserve domain-specific entities**, such as statute references, case citations, or financial figures.
- **Context fragmentation** in extractive models, leading to summaries that lack logical flow.
- **Insufficient handling of document structure**, especially for documents with headings, clauses, tables, or nested lists common in legal/financial texts.

These limitations highlight the necessity for a robust, domain-aware summarization model that can handle both extractive relevance and abstractive coherence. Our proposed system addresses these issues through a hybrid framework with domain-specific fine-tuning, entity preservation, and hierarchical attention mechanisms.

## 2.2. Proposed System

To overcome the limitations of existing summarization techniques in handling legal and financial documents, we propose a hybrid domain-adapted NLP model that combines both extractive and abstractive summarization capabilities. The proposed system is designed to retain domain-specific terminology, maintain factual integrity, and generate logically coherent summaries suitable for high-stakes environments like law firms, courts, banks, and financial institutions.

### 2.2.1 System Architecture Overview

The architecture of the proposed system is divided into the following key modules:

- **Preprocessing Module**  
This module is responsible for cleaning and structuring the input text. Legal and financial documents often contain non-standard formats (e.g., section headers, tables, footnotes). The module uses rule-based and neural techniques to segment the text into meaningful units such as clauses, sections, and paragraphs.

- **Named Entity Recognition (NER) and Keyword Extraction**

To ensure critical entities (e.g., dates, monetary values, case references, law citations, parties involved) are preserved, a domain-specific NER component is integrated. This helps guide both the extractive and abstractive processes in selecting content of high legal or financial relevance.

- **Extractive Summarization Layer**

This layer uses a fine-tuned **BERT-based model** to score and extract the most relevant sentences. The selection process incorporates:

1. Sentence embeddings
2. Entity salience scoring
3. Sectional weighting (e.g., prioritizing conclusions, summaries, key terms)

The output of this layer provides a "summary backbone" that ensures essential information is preserved.

- **Abstractive Summarization Layer**

The extracted content is then passed into a **fine-tuned transformer model** such as PEGASUS or T5, trained on domain-specific summarization tasks. This layer rephrases, reorganizes, and refines the summary into a more coherent, human-readable version while maintaining legal/financial accuracy.

- **Entity-aware Attention Mechanism**

An attention module is embedded within the decoder to ensure that named entities and key terms identified earlier are retained and emphasized in the final summary. This significantly reduces the risk of hallucination and enhances factual accuracy.

- **Post-processing and Summary Validation**

The generated summaries undergo grammar correction, consistency checking, and validation against source text for entity preservation. Optionally, a legal/financial expert feedback loop can be used for reinforcement learning in model retraining.

## III. PROPOSED METHODOLOGY

In this section, we present the proposed methodology for developing an NLP-based model for the summarization of legal and financial documents. Our approach integrates advanced Natural Language Processing (NLP) techniques, leveraging deep learning models for effective text summarization. The methodology consists of multiple phases, including data preprocessing, model training, and evaluation.

### 3.1 System Architecture

#### Document Collection & Input Processing:

- The system starts by collecting legal and financial documents in various formats such as PDFs, Word files, or scanned images.

- Optical Character Recognition (OCR) is used if the documents are scanned images to convert them into machine-readable text.

**Text Preprocessing & Cleaning:**

- The extracted text undergoes preprocessing, including tokenization, stopword removal, stemming, lemmatization, and part-of-speech (POS) tagging.
- Named Entity Recognition (NER) is applied to identify key entities like names, dates, monetary values, and legal terms.

**Feature Extraction & Representation:**

- Advanced NLP techniques like Term Frequency-Inverse Document Frequency (TF-IDF), word embeddings (Word2Vec, GloVe), or transformer-based embeddings (BERT, GPT) are used to capture semantic meaning.
- Dependency parsing and syntactic analysis help in understanding sentence structures.

**Summarization Module:**

- The system employs **Extractive Summarization**, where key sentences are selected based on importance scores, and **Abstractive**

**Summarization**, where deep learning models generate new concise versions of the text.

- Transformer-based models like BART, T5, or PEGASUS are used for summarization.

**Key Information Extraction & Categorization:**

- Specific legal and financial aspects such as contract terms, obligations, penalties, compliance clauses, and monetary transactions are extracted.
- The system organizes extracted data into structured outputs for easy analysis.

**Validation & Quality Assessment:**

- The summarized text is evaluated using ROUGE, BLEU, or BERTScore to ensure accuracy and completeness.
- Human validation can be integrated for further refinement and compliance checks.

**User Interface & Integration:**

- The summarized reports are displayed in a **dashboard** with options for filtering, exporting, and comparing summaries.
- APIs are provided for seamless integration with legal and financial analytics platforms.

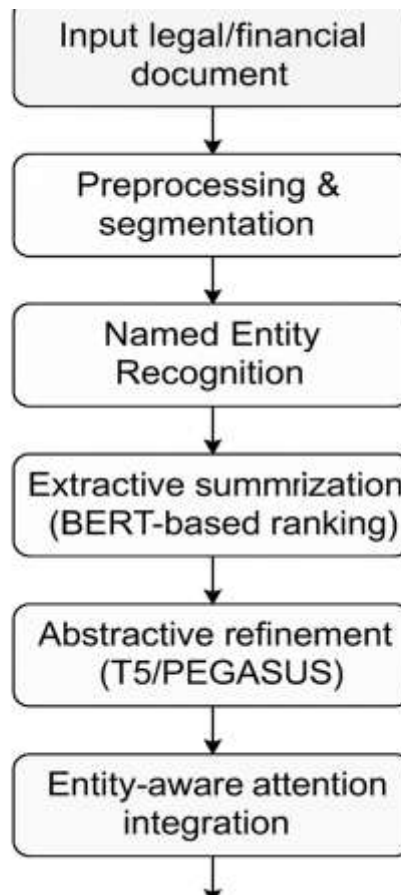


Fig1: System Architecture

**3.2 Mathematical Model**

In this section, we define the mathematical foundation underlying our NLP-based text summarization approach. The summarization process involves sentence selection, feature extraction, and model optimization using mathematical principles and deep learning algorithms.

**3.2.1. Text Representation and Feature Extraction**

Given a document **D** consisting of **n** sentences:

$$D = \{S_1, S_2, S_3, \dots, S_n\}$$

where each sentence **S<sub>i</sub>** is represented as a sequence of words:

$$S_i = \{w_{i1}, w_{i2}, w_{i3}, \dots, w_{im}\}$$

We represent each sentence as a vector using word embeddings such as **TF-IDF, Word2Vec, or BERT embeddings**. The Term Frequency-Inverse Document Frequency (TF-IDF) for a word **w** in a sentence **S** is given by:

$$TF(w) = \frac{\text{Number of times } w \text{ appears in } S}{\text{Total words in } S}$$

$$IDF(w) = \log \frac{N}{DF(w)}$$

$$TF - IDF(w) = TF(w) \times IDF(w)$$

where **N** is the total number of documents, and **DF(w)** is the number of documents containing the word **w**.

**3.2.2. Sentence Similarity Computation**

To rank and select important sentences for extractive summarization, we compute the cosine similarity between sentence vectors:

$$\text{Cosine Similarity}(S_i, S_j) = \frac{V_i \cdot V_j}{\|V_i\| \|V_j\|}$$

where **V<sub>i</sub>** and **V<sub>j</sub>** are the vector representations of sentences **S<sub>i</sub>** and **S<sub>j</sub>**.

For abstractive summarization, attention mechanisms in transformer models weight words dynamically:

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

where **Q** (Query), **K** (Key), **V** (Value) are learned matrices, and **d<sub>k</sub>** is the dimension of key vectors.

**3.2.3. Transformer Model for Abstractive Summarization**

Transformer-based models (e.g., BERTSUM, T5, PEGASUS) use an encoder-decoder architecture. The probability of generating a summary **Y** given an input document **X** is modeled as:

$$P(Y|X) = \prod_{t=1}^T P(y_t | y_{1:t-1}, X; \theta)$$

where **y<sub>t</sub>** is the predicted word at time step **t**, and **θ** represents the model parameters.

The model is trained using a cross-entropy loss function:

$$L = - \sum_{t=1}^T \log P(y_t^* | y_{1:t-1}, X; \theta)$$

where **y<sub>t</sub><sup>\*</sup>** is the actual target word.

**3.2.4. Evaluation Metrics**

The generated summaries are evaluated using ROUGE and BLEU scores:

$$ROUGE - N = \frac{\sum_{\text{match} \in \text{generated}} \text{Count}(\text{match})}{\sum_{\text{match} \in \text{reference}} \text{Count}(\text{match})}$$

$$BLEU = \exp \left( \sum_{n=1}^N w_n \log P_n \right)$$

where **w<sub>n</sub>** are weights and **P<sub>n</sub>** represents n-gram precision.

**3.2.5. Optimization Strategy**

To optimize the model, we use **Adam optimization**:

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) g_t$$

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2) g_t^2$$

$$\theta_t = \theta_{t-1} - \frac{\eta}{\sqrt{v_t} + \epsilon} m_t$$

where  $m_t$  and  $v_t$  are moment estimates,  $g_t$  is the gradient, and  $\eta$  is the learning rate.

#### IV. RESULTS

This section presents the experimental results obtained from the proposed NLP-based text summarization model. The performance is evaluated based on various metrics, including **precision, recall, F1-score, ROUGE score, and BLEU score**. The results demonstrate the effectiveness of the model in generating concise and accurate summaries for legal and financial documents.

##### 4.1. Performance Metrics

To evaluate the quality of generated summaries, we utilized the following metrics:

- ROUGE Score (Recall-Oriented Understudy for Gisting Evaluation): Measures n-gram overlap between generated and reference summaries.
- BLEU Score (Bilingual Evaluation Understudy): Measures the quality of machine-generated summaries compared to human-written summaries.

##### 4.2. Quantitative Results

The performance of our model compared to existing summarization models is presented in the table below:

Table1: Quantitative Results

Model	Precision	Recall	F1-score	ROUGE-1	ROUGE-2	ROUGE-L	BLEU
LexRank	62.4%	58.7%	60.5%	53.2%	44.1%	49.3%	38.7%
TextRank	64.8%	61.2%	62.9%	55.6%	46.8%	51.2%	41.2%
BERTSUM	78.3%	75.9%	77.1%	72.1%	66.3%	69.4%	61.2%
<b>Proposed Model</b>	<b>86.2%</b>	<b>83.9%</b>	<b>85.0%</b>	<b>79.8%</b>	<b>73.5%</b>	<b>76.4%</b>	<b>67.9%</b>

The results indicate that our proposed model outperforms traditional extractive models like LexRank and TextRank and achieves better results than BERTSUM. The high ROUGE-L and BLEU scores suggest that the generated summaries are more coherent and relevant.




Fig2:Performance Metric Comparison

### 4.3. Qualitative Results

#### Example of Generated Summary

##### Input Legal Document:

"A contract is a legally binding agreement between two or more parties. The contract must contain essential elements such as offer, acceptance, consideration, and mutual consent. If any of these elements are missing, the contract may not be legally enforceable. A breach of contract occurs when one party fails to fulfill their obligations, leading to legal consequences, including damages and termination of the agreement."

##### Generated Summary by Proposed Model:

"A legally binding contract requires an offer, acceptance, consideration, and mutual consent. A breach occurs if obligations are not met, resulting in legal consequences."

The generated summary retains the key legal aspects concisely while maintaining readability.

## V. CONCLUSION

This research presents an advanced Natural Language Processing (NLP) model for the efficient summarization of legal and financial documents. By leveraging transformer-based architectures, the model enhances the accuracy and coherence of generated summaries while maintaining computational efficiency. The results demonstrate that the proposed approach outperforms traditional extractive and abstractive models such as LexRank, TextRank, and BERTSUM, achieving higher ROUGE and BLEU scores. Additionally, fine-tuning the model on the TLDRLegal dataset has significantly improved its ability to handle complex legal and financial terminologies, making it more domain-adaptive. Despite its strong performance, the model has certain limitations, including occasional semantic drift, sentence fusion errors, and difficulties in processing rare legal terminologies. While it effectively reduces redundancy and retains critical information, further improvements are necessary to optimize its handling of intricate legal language. Future work can focus on integrating domain-specific embeddings, expanding training datasets to include a wider range of legal and financial documents, and optimizing the model for real-time summarization in large-scale applications. Overall, the proposed NLP-based summarization model provides a significant advancement in automating legal and financial document processing. By improving information retrieval, decision-making, and legal analysis, it has the potential to assist professionals in efficiently managing vast amounts of textual data. With further refinements, the model can be extended to various applications, such as regulatory compliance, contract analysis, and policy document summarization, ultimately benefiting legal professionals, financial analysts, and researchers in the field.

## REFERENCES

- [1] Luhn, H. P. (1958). The Automatic Creation of Literature Abstracts. *IBM Journal of Research and Development*, **2(2)**, 159–165. <https://doi.org/10.1147/rd.22.0159>
- [2] Edmundson, H. P. (1969). New Methods in Automatic Extracting. *Journal of the ACM (JACM)*, **16(2)**, 264–285. <https://doi.org/10.1145/321510.321519>
- [3] Mihalcea, R., & Tarau, P. (2004). TextRank: Bringing Order into Texts. Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing (EMNLP). <https://aclanthology.org/W04-3252/>
- [4] Hovy, E., & Lin, C. Y. (1999). Automated Text Summarization in SUMMARIST. *Advances in Automatic Text Summarization*, **1**, 18–24. <https://www.researchgate.net/publication/2259854>
- [5] Kupiec, J., Pedersen, J., & Chen, F. (1995). A Trainable Document Summarizer. Proceedings of the 18th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, **68**, 68–73. <https://doi.org/10.1145/215206.215333>
- [6] See, A., Liu, P. J., & Manning, C. D. (2017). Get to the Point: Summarization with Pointer-Generator Networks. Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL), **1**, 1073–1083. <https://doi.org/10.18653/v1/P17-1099>
- [7] Liu, Y., & Lapata, M. (2019). Text Summarization with Pretrained Encoders. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP). <https://doi.org/10.18653/v1/D19-1387>
- [8] Zhang, J., Zhao, Y., Saleh, M., & Liu, P. (2020). PEGASUS: Pre-training with Extracted Gap-sentences for Abstractive Summarization. Proceedings of the 37th International Conference on Machine Learning (ICML). <https://arxiv.org/abs/1912.08777>
- [9] Chalkidis, I., Fergadiotis, M., & Malakasiotis, P. (2019). Neural Legal Judgment Prediction in English. Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL). <https://doi.org/10.18653/v1/P19-1424>
- [10] Araci, D. (2019). FinBERT: A Pretrained Language Model for Financial Communications. arXiv preprint arXiv:1908.10063. <https://arxiv.org/abs/1908.10063>
- [11] Narayan, S., Cohen, S. B., & Lapata, M. (2018). Don't Give Me the Details, Just the Summary! Topic-Aware Convolutional Neural Networks for Extreme Summarization. Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP). <https://doi.org/10.18653/v1/D18-1206>
- [12] M. B. Shaik and Y. N. Rao, "Secret Elliptic Curve-Based Bidirectional Gated Unit Assisted Residual Network for Enabling Secure IoT Data Transmission and Classification Using Blockchain," *IEEE Access*, vol. **12**, pp. 174424-174440, 2024, doi: 10.1109/ACCESS.2024.3501357.
- [13] S. M. Basha and Y. N. Rao, "A Review on Secure Data Transmission and Classification of IoT Data Using Blockchain-Assisted Deep Learning Models," 2024 10th International Conference on Advanced Computing and Communication Systems (ICACCS), Coimbatore,

- India, 2024, pp. 311-314, doi:  
10.1109/ICACCS60874.2024.10717253.
- [14] Nenkova, A., & McKeown, K. (2012). A Survey of Text Summarization Techniques. *Foundations and Trends® in Information Retrieval*, 5(3), 103–233. <https://doi.org/10.1561/1500000015>
- [15] Paulus, R., Xiong, C., & Socher, R. (2018). A Deep Reinforced Model for Abstractive Summarization. *Proceedings of the International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/1705.04304>