Policy Document Summarization with Controlled Abstractive Techniques

by
Sehban Fazili

Under the supervision of
Dr. Md. Shad Akhtar

Submitted in partial fulfillment of the requirements for the degree of Master of Technology, CSE-AI

Department of Computer Science Engineering
Indraprastha Institute of Information Technology - Delhi
May, 2023

Certificate

This is to certify that the thesis titled *"Policy Document Summarization with Controlled Abstractive Techniques"* being submitted by **Sehban Fazili** to the Indraprastha Institute of Information Technology Delhi, for the award of the Master of Technology, is an original research work carried out by him under my supervision. In my opinion, the thesis has reached the standards fulfilling the requirements of the regulations relating to the degree.

The results contained in this thesis have not been submitted in part or full to any other university or institute for the award of any degree/diploma.

May,2023

Dr Md. Shad Akhtar

Department of Computer Science Engineering
Indraprastha Institute of Information Technology Delhi
New Delhi 110 020

Acknowledgements

I would like to express my deepest gratitude to my advisor Dr.Md Shad Akhtar for their invaluable guidance, support, and expertise throughout this journey. Their passion for teaching and commitment to excellence have been truly inspiring. Their mentorship has played a pivotal role in shaping my research skills and intellectual growth, and I am incredibly fortunate to have had the opportunity to learn from them.

I am also indebted to my lab mates for their camaraderie, collaboration, and stimulating discussions. Their diverse perspectives, unwavering support, and willingness to share their knowledge have enriched my research experience. I am grateful for the countless hours we spent working together, troubleshooting experiments, and brainstorming ideas. Their enthusiasm and dedication have been instrumental in creating a vibrant and motivating research environment.

Furthermore, I would like to extend my heartfelt appreciation to my friends outside the lab who have been a constant source of encouragement and motivation.

I would also like to acknowledge the broader research community for their contributions to the field. The wealth of knowledge and resources they have shared through publications, conferences, and online forums have been indispensable in shaping my understanding and pushing the boundaries of my research.

Lastly, I am immensely grateful to my family for their unwavering love, encouragement, and support. Their belief in me, even during challenging times, has been a constant source of strength.

Abstract

Privacy policies are often lengthy and complex, hindering individual's ability to make informed decisions about their data privacy. Abstractive summarization techniques can improve accessibility and transparency, but there is a lack of research in this area. Further development of these techniques can enhance comprehension of privacy policies and promote trust between individuals and organizations. In this work, we propose a controlled abstractive text summarization approach using a Bidirectional and Auto-Regressive Transformer (BART) model, which achieves state-of-the-art performance on our custom dataset. Our method optimizes the relevance and duration of the generated summaries to enable controlled summaries by integrating a reinforcement learning framework and a tailored loss function. We also introduce a new dataset of privacy policy documents and their summaries and establish performance benchmarks for future research. Experimental results on the custom dataset demonstrate significant improvement in summarization quality compared to several baseline methods, as measured by ROUGE and BLEU scores. The proposed approach has the potential to facilitate comprehension of privacy policies and improve user's privacy awareness.

Contents

List of Tables

List of Figures

List of Algorithms

Chapter 1

Introduction

Text Summarization is the process of condensing large amounts of information into a shorter form while retaining its most important content. The goal of summarization is to provide readers with a quick overview of the main points in a document or text without having to read the entire content. Summarization has become increasingly important due to the abundance of information available in today's digital age, and it is used in a variety of fields, including journalism, academic research, and business intelligence. There are two main approaches to summarization: extractive and abstractive. Extractive summarization involves selecting the most relevant sentences or phrases from a text and presenting them as a summary. This method is widely used in news articles and is based on statistical and linguistic algorithms that determine which sentences contain the most critical information. Abstractive summarization, on the other hand, involves generating a summary in natural language that captures the essence of the original text. This method uses machine learning algorithms such as neural networks and deep learning models to understand the context and meaning of the text and generate a summary.

1.1 Early Models

1.1.1 Latent Semantic Analysis (LSA)

Introduced in the 1990s by [4], LSA employed statistical methods to uncover the latent semantic structure of document collections. By representing documents and queries in a high-dimensional space, LSA aimed to capture the underlying relationships between words and sentences. LSA captures semantic relationships between words and documents by reducing the dimensionality of a term-document matrix through singular value decomposition (SVD).

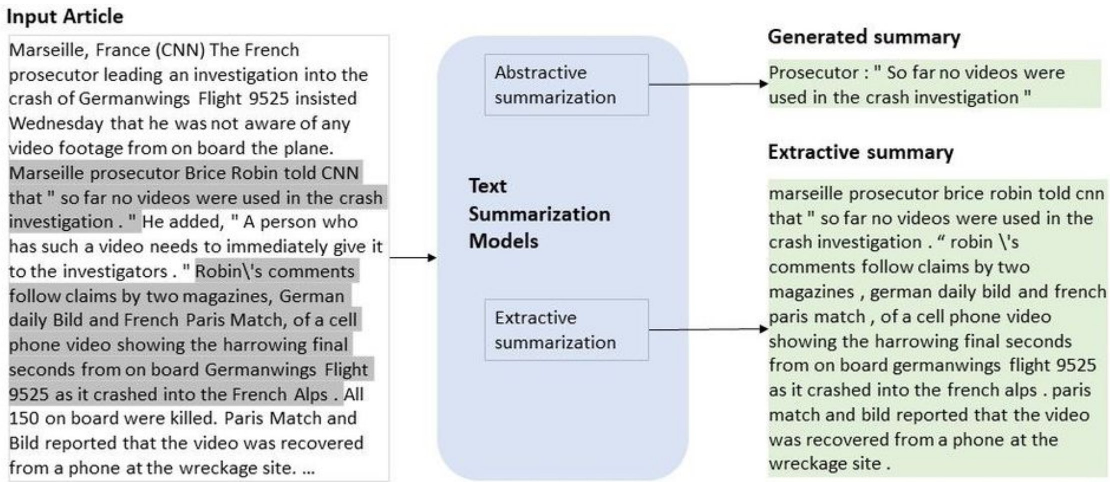

Figure 1.1: A comparison between extractive and abstractive summarization

LSA enables measuring word and document similarity, making it useful for tasks like information retrieval and document clustering. However, LSA has limitations, including its bag-of-words representation, lack of interpretability, sensitivity to noise, and challenges with sparse data.

1.1.2 LexRank

Proposed in 2004 by [6], LexRank was an influential graph-based algorithm for extractive summarization. It utilised a graph representation where sentences were vertices and edges represented sentence similarity. The application of a variant of the PageRank algorithm enabled the identification of important sentences for summary generation. It selects important sentences based on their similarity and centrality in the document. The algorithm effectively captures key information and can be applied to various languages. LexRank's strengths include content selection, language independence, and its graph-based representation. However, it may lack sentence cohesion, depend on the quality of similarity measures, be limited to extractive summarization, and face difficulties with specialised texts.

1.1.3 TextRank

Also introduced in 2004 by [18], TextRank was another significant graph-based algorithm for extractive summarization. Inspired by PageRank, TextRank assigned importance scores to sentences based on their centrality in the sentence graph. The most salient sentences were then selected to form the summary. It represents a text document as a graph, ranks the importance of words or sentences using iterative algorithms, and generates extractive summaries based on the top-ranked components. It is language-independent, scalable, and efficient for processing large documents. However, it may face challenges in maintaining overall coherence and generating abstractive summaries. Considerations for similarity measures and noisy data are important.

1.1.4 Multilingual Evaluation of Text Summarization Systems (MEAD)

MEAD, also developed in 2004 by [21], emerged as a comprehensive framework for extractive summarization. By integrating techniques like sentence clustering, scoring, and topic identification, it enabled the generation of summaries from multiple documents. MEAD played a pivotal role in evaluating and comparing various summarization approaches, offering standardized metrics and methodologies for fair system comparisons. However, MEAD's reliance on reference summaries and primary focus on extractive summarization may limit its coverage of all aspects of summarization quality.

Advancements and Current Trends The early models paved the way for subsequent advancements in text summarization. Notably, the advent of deep learning and neural networks has led to notable progress in abstractive summarization techniques. Models such as BERT and GPT have leveraged the power of these techniques, achieving state-of-the-art performance in generating abstractive summaries.

A privacy policy is a crucial component of any organization that collects, processes, and stores personal information. It is a legal document that outlines how an organization will handle personal data and how it will comply with applicable data protection laws and regulations. The purpose of a privacy policy is to provide transparency to individuals about how their personal information is being used and to establish trust with customers.

Privacy policies are essential for both businesses and individuals. For businesses, having a privacy policy can protect them from legal issues related to data privacy and help to establish trust with customers. For individuals, privacy policies provide transparency about how their personal information is being used and enable them to make informed decisions about whether to share their data with a particular organization.

Privacy policies are typically long and complicated, making it difficult for people to find the time to read them. Abstractive summarization techniques can potentially enhance accessibility and transparency by generating concise and coherent summaries of policy documents, which can make the content more comprehensible and manageable for the general public. Despite the potential benefits of abstractive summarization techniques for policy documents, there is still a lack of research specifically focused on summarizing privacy policy documents. Given the increasing concern over data privacy and security, it is crucial to ensure that individuals have a clear understanding of how their personal data is being collected, used, and protected by organizations. Summarizing privacy policy documents using abstractive techniques can provide individuals with an easily digestible summary of complex legal language and help to promote transparency and trust between organizations and their users. To address this, we propose an abstractive text summarization approach using a Bidirectional and Auto-Regressive Transformer (BART) model, which is a state-of-the-art pre-trained language model that has shown excellent performance on several natural language processing tasks. Our approach leverages the power of the BART model to generate high-quality summaries for privacy policy documents.

The main contributions of this paper are summarized as follows:

- We introduce a personalized loss function and a reinforcement learning framework to enhance the relevance and length of the produced summaries, ultimately resulting in the generation of controlled summaries.
- We introduce a new dataset of privacy policy documents and their summaries to evaluate the proposed approach. We manually annotate the summaries to ensure their quality and coherence.
- We also establish performance benchmarks for the proposed approach and several baseline methods on the custom dataset.

Experimental results demonstrate that our approach achieves state-of-the-art performance on the custom dataset, as measured by ROUGE and BLEU scores. Compared to several baseline methods, the proposed approach shows significant improvement in summarization quality. The proposed approach has the potential to facilitate comprehension of privacy policies and improve users' privacy awareness, as it generates concise and readable summaries that capture the key information of the original documents.

In the following sections we present the

1. Motivation
2. Related Work
3. Dataset
4. Methodology
5. Evaluation, and Results
6. Discussion and,
7. Future Scope

Chapter 2

Motivation

Privacy policies play a crucial role in explaining the procedures and guidelines governing the collection, utilization, and safeguarding of personal data by organizations. Nonetheless, these policies often exhibit repetition, complexity, and an abundance of legal terminology, rendering them difficult for individuals to comprehend in their entirety. This lack of understanding can lead to privacy concerns, as users might inadvertently provide consent for the sharing of sensitive information or remain unaware of their rights and alternatives.
The motivation behind privacy policy document summarization lies in addressing this problem by condensing lengthy and complicated policies into concise and easily comprehensible summaries. This research endeavor offers several key advantages:

- Accessibility: Summarized privacy policies serve as accessible resources, enabling individuals to conveniently access and comprehend relevant details regarding data collection, storage, and usage practices. Users can expeditiously grasp the fundamental tenets of a privacy policy without having to navigate through voluminous text, thereby augmenting transparency and empowering them to make well-informed decisions pertaining to their personal information.

- Time Efficiency: Privacy policies are known for their extensive length, and users often lack the time or patience to read them exhaustively. By providing short summaries, individuals can efficiently review and evaluate privacy practices, optimizing their time while still acquiring a clear comprehension of how their data will be handled.

- Empowerment and Control: Privacy policy summaries empower individuals by equipping them with the knowledge requisite for making privacy-conscious choices. Clear and concise summaries assist users in identifying potential privacy risks and making informed decisions regarding the disclosure of their personal information.

- Trust and Accountability: The process of summarizing privacy policies helps organizations be more transparent and build trust with their users. When organizations provide clear and easy-to-understand information, it shows that they are committed to being open and accountable. This strengthens their relationship with customers and promotes a culture of protecting privacy.

- Compliance: Privacy policy summarization also serves as a tool to aid organizations in adhering to legal and regulatory requirements. Summaries facilitate compliance with transparency and disclosure obligations, thereby enabling organizations to fulfill their legal responsibilities while upholding user trust.

Chapter 3

Related Work

3.1 Summarization with Pretrained Encoders:

Pretrained Encoders have been widely used for summarization in recent years. [16] proposed an unsupervised text summarization model based on BERT that achieved state-of-the-art results on several benchmark datasets. [26] explored various fine-tuning strategies for pretrained language models, including BERT and GPT, for summarization tasks, and found that using a small amount of labeled data for fine-tuning can significantly improve performance. [5] proposed a hierarchical transformer model for summarizing long documents, which achieved state-of-the-art results on several benchmark datasets. [27] proposed a new pre-training approach for summarization based on extracting gap sentences from a document, and training a transformer-based model to fill in the gaps, which resulted in a model called PEGASUS that achieved state-of-the-art results on several benchmark datasets. These recent works demonstrate the effectiveness of pretrained encoders for summarization tasks and highlight the importance of fine-tuning strategies and model architectures for achieving high performance.

3.2 Controlled Text Generation:

Controlled text generation is a task aimed at generating realistic sentences with desired attributes, such as sentiments or topics. Most efforts in controlled text generation rely on conditional pre-trained language models [3, 13]. [13] employ a GPT2-like pre-trained language model and train it from scratch using a large corpus that includes various control codes. Consequently, controlled generation is achieved by utilizing these control codes as prompting words. [3] propose a method that avoids additional training by combining the GPT-2 model with several simple attribute classifiers, whose gradients can update the latent representations. Another research direction explores the utilization of limited labeled data through learning latent representations [10]. [10] propose an approach to controlled text generation by learning disentangled latent representations, which include independent content and attribute components. In this paper, we learn and approach controlled text generation by decomposing the prior space into several parts. Another approach proposed in the literature combines an encoder-decoder architecture with a disentangled representation learning framework [12]. The proposed approach learns to generate text that reflects both the desired attributes and the style of the input text. The approach is evaluated on several datasets and shows promising results. Another approach by [11] proposes a method for controllable text generation that uses a pre-trained language model and a set of attribute classifiers to guide the generation process.

3.3 Reinforcement learning for summarization:

Reinforcement learning (RL) has been increasingly applied to the task of summarization in recent years. Several papers have proposed RL-based approaches for abstractive summarization that use a combination of supervised learning and RL, hierarchical RL-based methods that use a combination of word-level and sentence-level operations, and saliency-based attention mechanisms to improve the quality of the generated summaries [7, 20, 24, 25]. [23] introduced a RL-based approach to machine translation that uses simulated human feedback to train the model. Furthermore, [24] developed an RL-based approach for document summarization, which incorporates a saliency-based attention mechanism to enhance the overall quality of the produced summaries.[15] presented a RL-based framework for abstractive summarization that learns from human feedback, which can help overcome reward sparsity in RL. Lastly, [8] presented a general framework to train abstractive summarization models using question-answering based rewards in a reinforcement learning setting. The evaluation of the proposed framework encompasses three transformer-based summarization models, utilizing two publicly available datasets. The generated summaries are subjected to assessment through a combination of automatic evaluation measures and human judgments. The findings reveal that the integration of question-answering rewards within the framework yields promising outcomes, demonstrating its efficacy as a versatile approach for enhancing neural abstractive summarization.

Chapter 4

Dataset

The domain of Privacy Policy Summarization has not been extensively studied so far, and there is a lack of resources, such as datasets, that are specifically designed for this task. Therefore, we recognized the need for a novel dataset that is tailored to privacy policy documents and aims to facilitate research in this field. To create this dataset, we focused on privacy policies of websites, which are essential documents that outline how websites collect, use, and disclose personal information. Our dataset consists of a collection of privacy policy documents that are representative of various types of websites, such as e-commerce, social media, and news websites. Each privacy policy document in the dataset is annotated with three columns of information: the policy name, the original text, and the summary. The policy name column identifies the website and the specific privacy policy document being summarized. The original text column contains the full text of the privacy policy, while the summary column contains a concise and informative summary of the key points in the privacy policy.

The research project led by [1] aims to improve the accessibility and utility of privacy policies for various stakeholders in the web community, including researchers, users, regulators, and journalists. To accomplish this, the project involved curating and analyzing a massive dataset of one million privacy policies spanning multiple periods. The team used automated analysis methods to extract key information such as readability scores, topic models, textual similarity, and key phrases from the dataset. We utilized the curated privacy policies provided by [1] as our source text and generated summaries based on them.

Annotation Guidelines: To produce a summary for each privacy policy, we utilize the subsequent annotation guidelines:

- If multiple websites have identical policies, discard all but one.

- In case there are URLs linking to other websites, disregard them.

- Also, dismiss extremely brief documents that lack significant information.
- Refrain from including any policy content that is not relevant to the topic at hand.
- Also, dismiss the policies that are incomplete.

By following these guidelines, we aim to create a valuable summarization dataset that can be used for a wide range of research and practical applications.
Dataset Statistics: To assess our model's performance, we have split the dataset into two sets: a train set and a test set. The train set contains 1536 samples, while the test set has 384 samples. This split was done using an 80:20 ratio, which is a commonly used ratio for training and testing machine learning models. A brief statistic of the dataset is presented below:

Total Documents	1920
Total Paragraphs	1,20,991
avg_tokens_per_doc	1707.3
avg_tokens_per_summary	228.24

Table 4.1: Dataset Statistics

Figure 4.1 presents the word cloud visualization of the original text extracted from the dataset. This word cloud provides a graphical representation of the most frequently occurring words in the dataset's original text, with the size of each word indicating its relative frequency. By analyzing Figure 4.1, we can quickly identify the prominent terms and gain an understanding of the main themes and topics covered within the dataset's original text.
Similarly, Figure 4.2 depicts the word cloud of the summary column derived from the same dataset. The word cloud in Figure 4.2 highlights the most common words found in the summary column, offering insights into the key information and core concepts summarized in a condensed form.
Comparing the two word clouds, we can observe that the major keywords and important concepts remain consistent between the original text and the summary column. The similarity in the word clouds indicates that the essential content is preserved in the summary, ensuring that crucial details are not lost during the annotation process.
Our novel dataset is anticipated to serve as a valuable resource for researchers interested in the advancement and evaluation of automated techniques for summarizing privacy policies. By providing a comprehensive and diverse collection of privacy policy documents and their corresponding summaries, our dataset offers a solid foundation for conducting in-depth investigations and benchmarking the performance of summarization methods.

Figure 4.1: WordCloud for Original_Text

Figure 4.2: WordCloud for corresponding Summaries.

Chapter 5

Methodology

We follow the previous trends, which use the pre-trained model to help improve the task. As there has been little to no work on Privacy policy documents, traditional approaches perform well but have several problems. To solve the problems we face in applying these traditional approaches, we propose a novel model. Figure 5.1 provides an illustrative overview of our proposed model, comprising three distinct components: BART model, QA Generator, and the Reward model. The detailed explanation can be found in the next sub-sections.

5.1 Baseline Methods

To evaluate the effectiveness of the proposed methodology, we compare it against the following baseline approaches:

5.1.1 Extractive Oracle

An extractive oracle[9], is a system or model that leverages extractive techniques to provide summaries by extracting key information directly from the source text. The objective of an extractive oracle in text summarization is to identify the most important sentences or phrases within a document and assemble them into a concise summary. This is typically achieved through techniques such as sentence scoring or ranking, where sentences are assigned importance scores based on their relevance, informativeness, or other criteria. By utilising an extractive oracle, the summarization process becomes more efficient as it avoids the need for generating entirely new sentences or rewriting the content. Instead, it relies on the existing information within the document and extracts the most salient parts for the summary. However, it is crucial to acknowledge that extractive summarization has its limitations. Extractive oracles may struggle to capture the essence of the document as a whole or fail to convey the

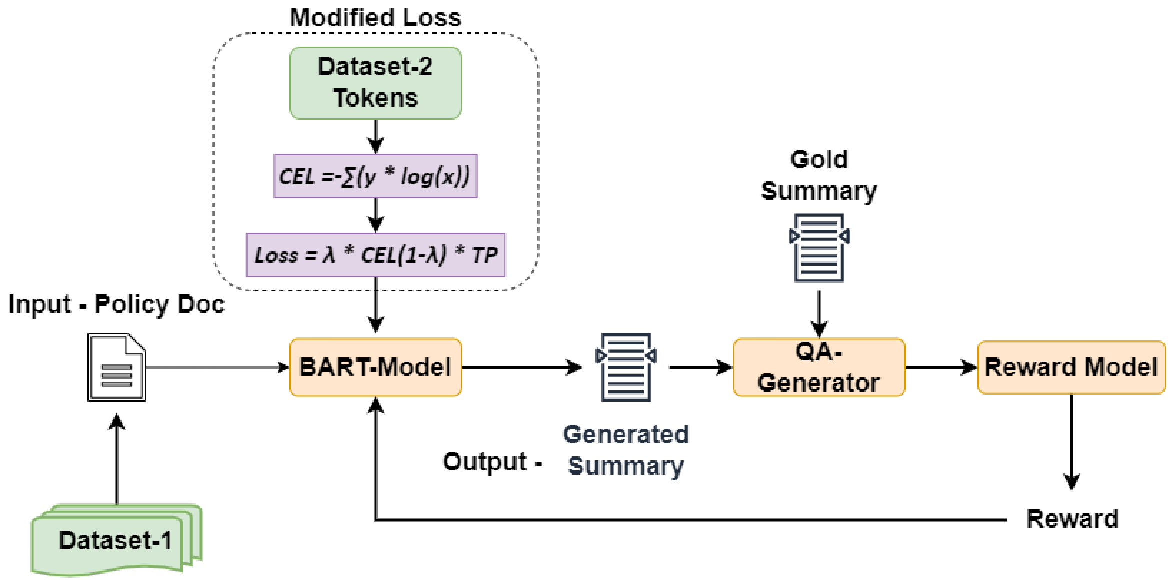

Figure 5.1: Model Architecture.

overall context and flow of the original text. They may also face challenges when dealing with documents that lack well-defined structures or contain intricate relationships between sentences. In the context of our study, we utilised this Extractive Oracle approach on our unique dataset. The process involves breaking down a paragraph of text into smaller sentences using the spaCy tool, which provides us with a list of sentences identified in the text. Next, we compared each sentence from the paragraphs to a target sentence and computed the ROUGE Score. This score quantifies the similarity between each sentence and the target sentence, enabling us to identify the most relevant sentence based on the highest score. This procedure was performed for all paragraphs, resulting in a collection of the most crucial sentences from each paragraph. The outcomes of this model can be observed in Table 2.

5.1.2 Bert2Bert

Bert2Bert revolves around utilising a BERT model for both encoding the source text and decoding the summary. This methodology capitalises on the strengths of BERT's pretraining on a vast amount of textual data, enabling it to capture the contextual information present in the source text and generate a coherent summary. In the context of summarization, Bert2Bert involves encoding the source text using the BERT model's encoder, which produces a representation of the input sequence. This representation is then passed to the BERT model's decoder, which generates the summary by predicting the next token based on the encoded context and previously generated tokens. By employing BERT as both the encoder and decoder, the model becomes capable of effectively comprehending the source text's semantics, relationships, and essential details. BERT's bidirectional nature enables it to capture contextual information from both preceding and succeeding words, thereby facilitating the generation of a cohesive and informative summary. To train the Bert2Bert model, we utilise a dataset consisting of paired source texts and their corresponding summaries, specifically curated from our novel Privacy Policy dataset. The results obtained can be seen in Table 2. The results highlight certain limitations, such as challenges in generating abstractive summaries, constraints on input length, difficulties in fine-tuning, and a potential tendency to overly emphasise the source text.

5.1.3 T5-Summarizer

The development of T5, a pre-trained language model by Google AI, marks a significant breakthrough in the realm of natural language processing (NLP). Its transformer-based architecture and unified text-to-text approach endow it with remarkable versatility, making it capable of undertaking various NLP tasks, including language translation, summarization, and question-answering. This unique architecture facilitates transfer learning, enabling the model to leverage its pre-trained knowledge to learn and generalize effectively on new tasks. Therefore, T5 has become a widely adopted and influential model within the NLP research community due to its effectiveness and applicability in various language-related tasks.
In the context of our study, we employ T5 as our baseline model and fine-tune it using our dataset. Fine-tuning allows us to adjust the model's pre-trained parameters, so it can better adapt to the specifics of our dataset and the tasks we aim to accomplish. By fine-tuning the T5 model, we aim to improve its performance and generate more accurate results in our experiments. The outcomes of this approach are presented in the following section.

5.1.4 BART-Summarizer

The Bidirectional and Auto-Regressive Transformer (BART) is a language model developed by Facebook AI Research (FAIR) based on the transformer architecture. The model has shown remarkable performance on a wide range of natural language processing (NLP) tasks, owing to its ability to perform both auto-regressive and bidirectional language modeling.
To evaluate its effectiveness, we fine-tuned the BART model on our dataset and report the results in the subsequent section. As a baseline model, BART allows us to assess the performance of our dataset against state-of-the-art results in NLP. By leveraging the strengths of BART's auto-regressive and bidirectional language modeling capabilities, we can gain insights into the nuances and intricacies of our dataset's linguistic structures.

5.1.5 PEGASUS

The PEGASUS model is a pre-trained transformer-based sequence-to-sequence architecture designed by Google AI, which has advanced the field of text summarization. One of its most notable features is the use of a novel pre-training objective known as ”gap-sentences generation”. This technique involves randomly masking input sentences, and training the model to predict the missing ones, improving the model's ability to understand the context and coherence of the text.
To evaluate the effectiveness of PEGASUS, we fine-tuned this model on our own dataset, and the results are presented in the next section.

5.2 Proposed Methodology

Abstractive summarization models are typically trained by minimizing the cross entropy loss of the reference summary at the word level. However, this approach does not inherently encourage models to prioritize factual accuracy with high precision and recall [17]. In order to address this limitation and enhance the factual accuracy of abstractive summarization, we have introduced a modified loss function. Moreover, we propose a general framework that leverages Question-Answering (QA)-based rewards and Reinforcement Learning (RL)-based training to further improve the summarization process. Figure 5.1 provides an illustration of our proposed framework, which encompasses the critical components involved. In the following sections, we outline these components in detail, emphasizing their significance in achieving our goal of enhancing factual accuracy in abstractive summarization.

5.2.1 BART Model with modified loss function for controlled summary generation

Using Dataset-1, which is our novel dataset we fine-tune the model and then change the loss function of the BART model that is we add some penalty to it based on the Dataset-2 which is our labs dataset based on NER task on the same privacy policy documents. We extract the tokens from the Dataset-2 and use that to modify the loss function of the BART model, that is, the Cross-Entropy Loss. Mathematically it can be seen as follows:

$$\text{CEL} = -\sum(y \cdot \log(x)) \tag{5.1}$$

$$\text{Loss} = \lambda \cdot \text{CEL} + (1 - \lambda) \cdot \text{TP} \tag{5.2}$$

Equation 5.1 represents the Cross-Entropy Loss (CEL), and Equation 5.2 represents the Loss function.
Where x is predicted class probabilities, and y is the true class probabilities. The overall loss function is a weighted sum of the cross-entropy loss and the forced token penalty, with lambda weight (λ) as the weighing factor. After performing this task, we get the controlled summaries.

5.2.2 Adding RL framework to improve Summaries

To improve the quality of summaries, we put forth a comprehensive framework that uses the power of QA-based rewards and RL-based training. By integrating QA generators as rewards, our proposed framework aligns with the principles of human evaluation, as it assesses the summary's relevance and accuracy based on its ability to effectively address questions originating from the source text. An added advantage of training with QA-based rewards is the encouragement of informativeness. QA models are trained to provide precise and informative answers, and incorporating them as rewards compels the RL agent to generate summaries that not only capture key information but also deliver accurate and meaningful responses. Furthermore, the utilisation of QA-based rewards stimulates the generation of multi-faceted summaries. The RL agent is incentivized to cover diverse aspects of the source text in order to effectively respond to a wide range of questions. This aspect ensures that the resulting summaries are comprehensive, encompassing different angles and dimensions of the original content. Figure 5.1 illustrates our proposed framework. Further elaboration on the Question-Answer Generator and reward model is provided in the subsequent sections.
Question-Answer Generator We use two pre-trained language models: one to generate questions based on the reference summary and another to answer those questions using both the generated and reference summaries. To calculate the reward score, we calculate the overlap between the generated and reference summaries using the metric called Rouge. If the Rouge score is higher than a specified threshold, the generated summary is considered good and is given a positive reward score. If the Rouge score is lower than the threshold, the generated summary is not considered good and is given a reward score of zero. Algorithm 1 provides a detailed explanation of the operational framework of the QA model. For the generation of questions from a summary, we utilize the pre-trained model "facebook/bart-large-cnn" developed by Facebook AI Research. This model has undergone fine-tuning on the CNN/Daily Mail dataset, which comprises news articles and their respective summaries. To determine the answer for a generated question derived from a summary, we employ an extractive QA model trained specifically on the SQuAD (Stanford Question Answering Dataset) task. By leveraging this model, we are able to identify and extract relevant answers from the given summary based on the generated questions. [22].

Algorithm 1 Rewards Calculation QA Framework

Input: Generated Summary (S), Question Generation Model (QG), Question Answer Generation Model (QA), Rouge Threshold (RT)
Output: The value of the reward (R) assigned to the generated summary (S).

1. Generate reference questions from the reference summary using QG
2. Initialize empty lists for generated answers and reference answers
3. For each question in reference_questions:
   - Generate an answer using the QA for the generated question and the original summary
   - Append the generated answer to the generated answers list
   - Generate an answer using the QA for the generated question and the reference summary
   - Append the reference answer to the reference answers list
4. Score = 0
5. Create a Rouge scorer object with the specified Rouge metric
6. For each generated answer and reference answer:
   - Calculate the Rouge scores between the generated answer and reference answer
   - Add the Rouge F-measure score to the Score
7. Calculate the average Rouge score
8. If the average Rouge score > RT:

   $$R = \text{mean(Rouge score)}$$

   Else: R = 0

10. Return R

Reward Model The reward function employed in this study is based on the semantic similarity between the answers extracted from the generated and ground truth summaries. Specifically, a generated summary is deemed relevant if it can accurately answer the questions extracted from the ground truth summary, indicating that it contains the essential information. On the other hand, a generated summary is considered factual if it can be accurately queried by the ground truth summary, i.e., the answers to the questions generated from the generated summary are consistent with those derived from the ground truth summary. The reward function is utilized by the reinforcement learning (RL) framework, as shown in Figure 5.1, to fine-tune the summary generation model. The Q-learning algorithm is employed to update the model's parameters in response to the rewards received from the environment. The model is trained to maximize the expected cumulative reward over time, leading to the generation of informative and coherent summaries. By leveraging the semantic similarity between the generated and ground truth summaries, the proposed reward function provides a more informative and effective feedback signal for training the summary generation model. Furthermore, the RL framework enables the model to learn from its mistakes and improve its performance through iterative interactions with the environment.

Chapter 6

Evaluation, and Results

In this chapter, we provide further details regarding the conducted experiments, the results of our evaluations, and the necessary analyses conducted as components of the results.

6.1 Evaluation Metrics

The experiments were conducted using the specified train and test splits. To assess the quality and effectiveness of the generated outputs, we employed a set of established evaluation metrics, including BLEU-1, BLEU-2, BLEU-3, and BLEU-4 [19], as well as ROUGE-1, ROUGE-2, ROUGE-L [14], and METEOR [2]. These metrics were selected due to their well-established usage in the field and their ability to provide a comprehensive assessment of the generated summaries. Each metric offers a unique perspective on the effectiveness and accuracy of the generated results, enabling a thorough analysis and comparison of the model's performance. By utilizing these multiple evaluation metrics, we were able to obtain a comprehensive understanding of the model's capabilities and make informed decisions based on the obtained results. Below, we provide a more detailed explanation of the evaluation metrics used in our study:

- ROUGE (Recall-Oriented Understudy for Gisting Evaluation) is a widely used evaluation metric for assessing the quality of text summarization. Its focus on recall measures how effectively the candidate summary captures information from the reference summaries. ROUGE encompasses various variants, including ROUGE-N, ROUGE-L, ROUGE-W, and ROUGE-S, each employing different techniques to calculate recall. In our task, we use ROUGE-N and ROUGE-L.

    1. ROUGE-N: This variant evaluates recall by considering n-grams, which are contiguous sequences of words. For example, ROUGE-1 measures the match be-

tween individual words (unigrams), while ROUGE-2 evaluates the match between adjacent word pairs (bigrams), and so on. For a given value of n, we count the total number of n-grams across all the reference summaries. Then, we determine how many of these n-grams are present in the candidate summary. The metric value is obtained by calculating the fraction of matching n-grams over the total number of n-grams. In our task, we employ ROUGE-1 and ROUGE-2.

2. ROUGE-L: This variant goes beyond simple recall and incorporates precision and recall into a single metric called the F-score. ROUGE-L uses the concept of the longest common subsequence (LCS) between the candidate and reference summaries. Given a candidate summary A of length m and a reference summary B of length n, we calculate precision (P) and recall (R) as follows:

$$\begin{aligned} \text{Precision (P)} &= \frac{\text{LCS}(A, B)}{m} \\ \text{Recall (R)} &= \frac{\text{LCS}(A, B)}{n} \\ \text{F1-score (F)} &= \frac{(1+\beta^2)\cdot(P\cdot R)}{\beta^2\cdot P + R} \end{aligned} \tag{6.1}$$

- BLEU (Bilingual Evaluation Understudy) score is a popular metric for evaluating the quality of machine-generated translations or summaries. It assesses the similarity between the generated output and reference translations or summaries. The BLEU score takes into account n-grams, which are contiguous word sequences, to measure precision and consider exact match. It combines the precision of matching n-grams with a penalty for shorter outputs. Mathematically it is given as:

$$\text{BLEU} = \text{BP} \cdot \exp\left(\sum_{i=1}^{n} w_i \cdot \log(p_i)\right) \tag{6.2}$$

where: BP is the brevity penalty factor,w_i is the weight assigned to each n-gram precision,p_i is the n-gram precision.

$$\text{BP} = \min\left(1, \exp\left(1 - \frac{r}{c}\right)\right) \tag{6.3}$$

where: r is the effective reference length,c is the length of the generated output.

- METEOR (Metric for Evaluation of Translation with Explicit Ordering) is an evaluation metric used to assess the quality of machine translation. It considers precision and recall at the unigram level. METEOR aligns candidate and reference translations at the phrase level, penalizes overgeneration and undergeneration, and calculates the F-measure as the final score. It is known for its linguistic robustness and comprehensive evaluation approach.

The score generated by these evaluation metrics ranges from 0 (worst) to 1 (best).

6.2 Implementation Details

We conducted experiments on the Noval Dataset, which we partitioned into a train-test ratio of 80:20. The train set consisted of 1536 samples, while the test set comprised 385 samples. For our experiments, we utilized the BART model as the base model. To enhance the summarization process, we introduced modifications to the loss function. Specifically, we replaced the cross-entropy loss function with a novel loss function, as explained in Section 4 of our study. Moreover, to enhance the quality of the generated summaries, we integrated a reinforcement learning (RL) framework into our approach. Within the RL framework, we employed the Q-Learning algorithm and fine-tuned the hyperparameters. The following values were set after thorough experimentation:

- Alpha: 0.1
- Gamma: 0.9
- Epsilon: 0.1
- Maximum Steps: 5
- Maximum Iterations: 50

6.3 Results

Given the novelty of the undertaken problem, we incorporated several existing systems to facilitate a comprehensive comparison. In order to establish a benchmark for evaluation, we adapted various related systems that have been previously developed and tested. By leveraging these existing systems, we were able to gauge the performance and effectiveness of our proposed approach in relation to established methods in the field. This comparative analysis provided valuable insights and allowed us to assess the novelty and potential advantages of our solution. Table 6.1 shows the experimental results for the baselines.

Table 6.1: Baseline Results

Model	Rouge			BLEU				METEOR
	R1	R2	RL	B1	B2	B3	B4	
Extractive Oracle	0.43	0.30	0.42	0.14	0.12	0.11	0.10	0.25
Bert2Bert	0.25	0.04	0.22	0.22	0.08	0.10	0.16	0.25
T5-Summarizer	0.44	0.24	0.42	0.32	0.24	0.19	0.17	0.34
BART	0.46	0.29	0.44	0.31	0.25	0.22	0.20	0.33
PEGASUS	0.35	0.17	0.32	0.18	0.12	0.09	0.08	0.23

6.3.1 Loss Results

We enhance the baselines specifically T5, BART and PEGASUS by introducing a modified loss function that evaluates the dissimilarity between the predicted and actual probability distributions. Our comprehensive loss function combines the cross-entropy loss, which serves as the original loss function, with the inclusion of a forced token penalty. To achieve this, we introduce a weighting factor, denoted as lambda (λ), to adjust the relative importance of these components. For a more detailed explanation, please refer to Chapter 5.
The experimental results for the baselines are presented in Table 6.2, highlighting the effectiveness of our modifications. Notably, the results demonstrate significant improvements in performance.

Table 6.2: Baseline Results with addition of modified Loss

Model	Rouge			BLEU				METEOR
	R1	R2	RL	B1	B2	B3	B4	
T5-Loss	0.45	0.26	0.43	0.32	0.24	0.19	0.17	0.34
BART-Loss	0.48	0.31	0.46	0.31	0.25	0.22	0.20	0.34
PEGASUS-Loss	0.38	0.21	0.36	0.23	0.17	0.14	0.12	0.26

6.3.2 Main Results

Table 6.3 presents the results of the proposed methodology. The findings demonstrate the performance of our model, CPD-Sum (Controlled Policy Document Summarization). A thorough analysis of the table reveals noteworthy improvements when compared to the baselines. Specifically, there is a remarkable increase in the R1 value by 0.14, indicating a substantial enhancement in the model's ability to capture important information from policy documents. Moreover, the R2 value shows a notable improvement of 0.5, suggesting a significant boost in the model's capability to generate concise and coherent summaries. Additionally, the RL value shows an increase of 0.2, indicating enhanced overall summarization quality. Although the remaining values in the table exhibit minimal variations, these results reaffirm the superiority of our CPD-Sum model in comparison to the baselines. Such improvements underscore the effectiveness and potential of our proposed methodology in the field of controlled policy document summarization.

Table 6.3: Results of the proposed methodology

Model	**Rouge**			**BLEU**				**METEOR**
	R1	**R2**	**RL**	**B1**	**B2**	**B3**	**B4**	
CPD-Sum	0.50	0.34	0.46	0.27	0.22	0.20	0.19	0.31

Chapter 7

Conclusion

In this work, we presented a novel approach for abstractive summarization of privacy policy documents utilizing the power of a Bidirectional and Auto-Regressive Transformer (BART) model. Our approach aimed to address the challenge of generating controlled and informative summaries that capture the essence of complex privacy policies. To achieve this, we introduced a customized loss function and incorporated a reinforcement learning framework, enabling us to optimize both the relevance and length of the generated summaries. To facilitate the evaluation and advancement of research in this domain, we also introduced a new dataset comprising privacy policy documents and their corresponding summaries. This dataset serves as a valuable resource for future investigations and provides a benchmark for assessing the performance of different summarization models. The experimental results obtained from our comprehensive evaluations highlight the effectiveness of the proposed approach. Our model achieved state-of-the-art performance on the custom dataset. The controlled generation of summaries allows for improved accessibility and transparency for users, enabling them to quickly grasp the key points of privacy policies without getting overwhelmed by excessive information. The findings of our work demonstrate the potential of our approach to make a significant impact in the field of privacy policy summarization. By addressing the critical need for concise and user-friendly representations of privacy policies, we contribute to enhancing user understanding. The implications of our work extend to various domains where privacy policies play a crucial role, including data protection, online services, and legal compliance.

Chapter 8

Future Scope

8.1 Incorporating User Feedback

To enhance the usability and usefulness of summarization tasks, it is important to incorporate user feedback. User feedback can provide valuable insights into the quality and relevance of the generated summaries, enabling iterative improvements. By considering user feedback, summarization systems can be tailored to meet individual needs, improving user satisfaction and overall utility.

8.2 Personalised Privacy Policy Summaries

Privacy preferences can vary among individuals, and tailoring privacy policy summaries to the specific needs and concerns of users can greatly enhance their understanding. Future developments may involve incorporating personalisation techniques that take into account user preferences, such as highlighting sections that are most relevant to their interests or providing summaries that align with their privacy preferences. This level of customization can empower users to easily grasp the privacy implications that matter most to them.

8.3 Multilingual Privacy Policy Summarization

With the expansion of online services globally, privacy policy summarization needs to address multilingual challenges. Future directions can focus on developing techniques to automatically summarise privacy policies in multiple languages. This would facilitate comprehension for users who are not proficient in the language of the original policy, helping them understand the privacy implications in their native language.

8.4 Compliance Verification and Transparency

Privacy policy summarization can also be utilised to ensure compliance with privacy regulations. Future advancements may involve developing algorithms that can analyse and summarise privacy policies in a way that enables regulatory authorities and organisations to verify compliance easily. These summaries can provide transparency and aid in audits, making it easier to identify potential violations and discrepancies between policies and actual data practices.

8.5 Explainable Privacy Policy Summaries

Privacy policy summarization should not only provide condensed versions of policies but also explain the rationale and implications behind various clauses and practices. Future directions can focus on generating explainable summaries that highlight the reasons behind data collection, sharing, and usage practices outlined in privacy policies. This can foster transparency and trust between users and organisations, enabling users to make more informed decisions about their data.

Bibliography

[1] Ryan Amos and et al. Privacy policies over time: Curation and analysis of a million-document dataset. In *Proceedings of the Web Conference 2021*, pages 1–12. ACM, 2021.

[2] Satanjeev Banerjee and Alon Lavie. Meteor: an automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005.

[3] Ramesh Dathathri, Andrea Madotto, and Zihang Lan. Plug and play language models: A simple approach to controlled text generation. *arXiv preprint arXiv:1912.02164*, 2020.

[4] Scott Deerwester, Susan T Dumais, George W Furnas, Thomas K Landauer, and Richard A Harshman. Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6):391–407, 1990.

[5] Li Dong, Furu Wei, Ming Zhou, and Ke Xu. Hierarchical transformers for long document summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics*, 2021.

[6] Günes Erkan and Dragomir R. Radev. Lexrank: Graph-based lexical centrality as salience in text summarization. *Journal of Artificial Intelligence Research*, 22:457–479, 2004.

[7] Jiatao Gu and et al. Deep reinforcement learning for sequence-to-sequence models. *arXiv preprint arXiv:1609.06018*, 2016.

[8] Chamara Gunasekara and et al. Using question answering rewards to improve abstractive summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 518–526, 2021.

[9] Tsutomu Hirao, Masaaki Nishino, Jun Suzuki, and Masaaki Nagata. Enumeration of extractive oracle summaries. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 386–396, Valencia, Spain, April 2017. Association for Computational Linguistics.

[10] Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P Xing. Toward controlled generation of text. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1587–1596, 2017.

[11] Haoyue Huang, Xiaochuan Wang, and Jing Zhu. Revisiting controllable text generation via plug-and-play language models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL'22)*, pages 5366–5377, 2022.

[12] Wenxuan Jin, Chenglin Tan, and Yichuan Zhou. Towards style-controllable and attribute-controllable text generation with disentangled representations. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 113–123, 2022.

[13] Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, and Caiming Xiong. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*, 2019.

[14] Chin-Yew Lin. Rouge: a package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.

[15] Fei Liu and et al. Learning to summarize from human feedback. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020.

[16] Yang Liu, Ming Li, Xin Liu, Yao Chen, and Maosong Sun. Presumm: A bert-based unsupervised text summarization model. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020.

[17] Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. On faithfulness and factuality in abstractive summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, 2020.

[18] Rada Mihalcea and Paul Tarau. Textrank: Bringing order into texts. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 404–411, 2004.

[19] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics (ACL)*, volume 1, pages 311–318, 2002.

[20] Romain Paulus and et al. A deep reinforced model for abstractive summarization. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2017.

[21] Dragomir R. Radev, Hongyan Jing, M. Stys, and David Tam. Centroid-based summarization of multiple documents: sentence extraction, utility-based evaluation, and user studies. In *Proceedings of the 2004 Document Understanding Conference (DUC)*, 2004.

[22] Pranav Rajpurkar, Robin Jia, and Percy Liang. Know what you don't know: Unanswerable questions for squad. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, 2018.

[23] Mathieu-Auguste Rondeau and et al. Reinforcement learning for bandit neural machine translation with simulated human feedback. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018.

[24] Xiaojun Wan and et al. Improving abstractive document summarization with salient information modeling. In *Proceedings of the 27th International Conference on Computational Linguistics*, 2018.

[25] Chenguang Wang and et al. A hierarchical reinforced sequence operation method for abstractive summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018.

[26] Liyuan Wang, Yue Zhang, Fuli Xu, and Xu Chen. Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping. *arXiv preprint arXiv:2103.05447*, 2021.

[27] Jingqing Zhang, Yao Zhao, Mohamed Saleh, and Peter Liu. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, 2020.

Appendix A

Appendix

A.1 Qualitative Analysis

The generated summaries produced by various models for comparative analysis with the gold summaries are presented below.

BART
Generated Summary TechBuzz.net, like most other websites, collects and uses the data contained in log files. Your IP address, your ISP (internet service provider, such as AOL or Shaw Cable), the browser you used to visit our site (such as Internet Explorer or Firefox), the time you visited our site, and which pages you visited throughout our site are all stored in log file data. Cookies are used by Google, a third-party vendor, to serve ads on your site. Third-party advertisements are also used by TechBuzz. Net to help fund the site. When they advertise on our site they may use technology such as cookies and web beacons, which will also send these advertisers (like Google) information such as your IP address. Third parties may use cookies and Web beacons to deliver adverts on our website.', ”for each visit to our website, our web server automatically recognizes only the visitor's IP address, but not the email address unless the visitor volunteers their email address to us by filling out a web form. We electronically collect the email addresses of those who complete online web forms and of those that communicate with us via email. We aggregate information on what pages and other resources users access or visit, as well as users' specific information. **Gold Summary**

continued from previous page

BART

We at Tech-Buzz.net understand how vital it is for you to keep your personal information private. Your personal information is never sold to a third party. Tech-Buzz.net gathers and uses log files in the same way that most other websites do. Your IP address, your ISP, the browser you used to visit our site, the time you visited the site, and which pages you visited across the site are all stored in log files. When you visit the site, Tech-Buzz.net uses cookies to record information such as your personal preferences. Cookies are used by Google, a third-party provider, to deliver adverts on your site. On Tech-Buzz.net, we also use third-party adverts to help fund the site. When they advertise on the site, some of these advertisers may use technology like cookies and web beacons, and they may also share the information they collect with others.

BART-Loss

Generated Summary

Yatra respects your privacy and recognizes the need to protect the personally identifiable information (any information by which you can be identified, such as name, gender, email address, postal address, frequent fly number, age, telephone number, etc.) you share with us. By using or accessing the Website or other Sales Channels, the User hereby agrees with the terms of this Privacy Policy and the contents herein. This Privacy Policy does not apply to any website(s), mobile sites and mobile apps of third parties, even if their websites/products are linked to our Website

Gold Summary

continued from previous page

BART
Yatra respects your privacy and recognizes the need to protect the personally identifiable information (any information by which you can be identified) you share with us. This Privacy Policy does not apply to any website(s), mobile sites and mobile apps of third parties, even if their websites/products are linked to our Website. With specific reference to booking/e-commerce transactions, Yatra collects the following personal sensitive information from you while transacting through Yatra: Name and Gender Phone Number, Address, Credit Card details, Date of birth in case of a child, Passport Number & frequent flyer number. Yatra does not sell or trade upon any of the above foregoing information without the consent of the user or customer. Some of our web pages use "cookies" so that we can better serve you with customised information when you return to our site. The information we collect and analyse is used to improve our service to you. We use third-party service providers to serve ads on our behalf across the Internet and sometimes on this site. They may collect anonymous information about your visits to our website, and your interaction with our products and services. This anonymous information is collected through the use of a \'pixel tag\'. Yatra takes appropriate steps to protect the information you share with us. We may, in good faith, disclose your personal sensitive information to protect ourselves from any liability or fraud or similar situation. The user\'s personal information may be disclosed for regulatory/ internal compliance and audits. We will not access your photos or camera without first getting your permission and we will never scan or import your photo library or camera roll. We will not scan or import your contacts stored on your phone without first getting your explicit permission. when you give us permission it enables us to put your travel plan on your calendar. We will not gather or use the specific location of your mobile device (by using, for example, GPS or Bluetooth) without first getting your explicit permission. This application reads, uploads and syntactically analyses text messages from travel providers within the app for any flight and train PNR booking messages. The User allow the Owner to provide its services to third party services' accounts for Displaying content from external platforms, SPAM protection, Hosting and backend infrastructure, Infrastructure monitoring, Content commenting, Interaction with external social networks and platforms, Location-based interactions, Social features, Contacting the User, User database management, Targeted Advertising and Remarketing and behavioural targeting.
CPD-Sum
Generated Summary

continued from previous page

BART

Digital Hill may disclose personally identifiable information under special circumstances, such as to comply with subpoenas or when your actions violate the Terms of Service. You are required to log-in to your Facebook TabSite after a certain period of time has elapsed to protect you against others accidentally accessing your account contents.

Gold Summary

Digital Hill collects the e-mail addresses of those who communicate with us via e- mail, aggregate information on what pages consumers access or visit, and information volunteered by the consumer (such as survey information and/or site registrations). The information we collect is used to improve the content of our Web pages and the quality of our service, and is not shared with or sold to other organisations for commercial purposes, except to provide products or services you've requested, or, under the following circumstances: We transfer information about you if Digital Hill or Facebook Tabsite is acquired by or merged with another company. A cookie is a small amount of data, which often includes an anonymous unique identifier, that is sent to your browser from a web site's computers and stored on your computer's hard drive. Digital Hill uses third party vendors and hosting partners to provide the necessary hardware, software,, networking, storage, and related technology required to run Facebook TabSite. Facebook TabSite or Digital Hill may disclose personally identifiable information, under special circumstances, such as to comply with subpoenas or when your actions violate the Terms of Service.

The table presented above provides a comprehensive analysis comparing the generated summaries with the gold standard summaries from different models. Our focus was primarily on the top-performing models, which included BART, BART-Loss, and CPD-Sum (our proposed model). By conducting this comparison, we aimed to identify any disparities between the generated summaries and the reference summaries.

After carefully examining the summaries, we noticed some important things. First, all the summaries were short and to the point, giving a condensed version of the original text. However, the BART model had fewer important details compared to the summaries generated by the BART-Loss and CPD-Sum models. While all three models exhibited good performance in terms of grammatical accuracy and overall fluency, the CPD-Sum model demonstrated superior language proficiency. Its summaries boasted a higher level of grammatical precision, coherently presenting the essential details with greater finesse and coherence. These findings help us understand the subtle differences and unique features of the models we analyzed. The BART-Loss and CPD-Sum models outperformed the BART model in terms of information retention, ensuring a more comprehensive summary. Additionally, the CPD-Sum model stood out as the most linguistically proficient, crafting summaries that excelled in both grammatical structure and overall readability.

A.2 Dataset Sample

Within this section, we present a sample of the dataset. The dataset comprises three columns, namely Policy_Names, Original_Text, and Summary, each serving a specific purpose in capturing essential information. The first column, Policy_Names, records the names or labels associated with the policies under consideration. These names provide a contextual reference to the particular policies within the dataset.
The second column, Original_Text, contains the original text content of the policies. This column serves as a repository of the comprehensive and detailed information present in the policies, offering a textual representation of their content, clauses, and provisions. It acts as a valuable resource for understanding the policies in their entirety.
The third column, Summary, is the summary of the policy or Original_Text.

	Policy_Names	Original_Text	Summary
0	007names.com	# Privacy Policy 007Names, Inc. knows that you...	We collect and store any information entered o...
1	01.com (zmailcloud)	### ZMAILCLOUD WEBSITE PRIVACY POLICY The Inte...	Without your permission, we will not sell your...
2	011networks.com	# Privacy Policy for 011 Networks 1) **Privacy...	This Online Privacy Policy may be modified at ...
3	012global.com	# 012Global | Low-Cost International Calls 1. ...	012global collects information from you that m...
4	01webdirectory.com	# 01Webdirectory.com: Privacy Policy 01webdire...	01webdirectory.com may place a "cookie" on you...
...	...	...	...
1916	Zkoss.org	Thanks for entrusting Potix Corporation, who c...	We collect information such as your browser ty...
1917	Zkteco.com	**Privacy Protection Policy** **Important Noti...	ZKTeco collects your personal and non-personal...
1918	Zkzi.com	**Privacy Policy for www.zkzi.com** If you req...	www.zkzi.com, like many other websites, makes ...
1919	Zoknowsgaming.com	**What information do we collect?** We collect...	When you register on our site, we collect info...
1920	Ztod.com	This Privacy Policy applies to the following w...	Ztod.com is devoted to safeguarding your perso...

1921 rows × 3 columns

Figure A.1: Dataset