



A Systematic Survey of Text Summarization: From Statistical Methods to Large Language Models

HAOPENG ZHANG, Department of Information and Computer Sciences, University of Hawaii at Manoa, Honolulu, United States

PHILIP S. YU, University of Illinois at Chicago, Chicago, United States

JIawei ZHANG, IFM Lab, Department of Computer Science, University of California, Davis, Davis, United States

Text summarization research has undergone several significant transformations with the advent of deep neural networks, pre-trained language models (PLMs), and recent large language models (LLMs). This survey thus provides a comprehensive review of the research progress and evolution in text summarization through the lens of these paradigm shifts. It is organized into two main parts: (1) a detailed overview of datasets, evaluation metrics, and summarization methods before the LLM era, encompassing traditional statistical methods, deep learning approaches, and PLM fine-tuning techniques, and (2) the first detailed examination of recent advancements in benchmarking, modeling, and evaluating summarization in the LLM era. By synthesizing existing literature and presenting a cohesive overview, this survey also discusses research trends, open challenges, and proposes promising research directions in summarization, aiming to guide researchers through the evolving landscape of summarization research.

CCS Concepts: • **Computing methodologies** → **Natural language generation**;

Additional Key Words and Phrases: Summarization, large language model, deep learning, dataset

ACM Reference Format:

Haopeng Zhang, Philip S. Yu, and Jiawei Zhang. 2025. A Systematic Survey of Text Summarization: From Statistical Methods to Large Language Models. *ACM Comput. Surv.* 57, 11, Article 277 (June 2025), 41 pages. <https://doi.org/10.1145/3731445>

1 Introduction

Text summarization is one of the most critical and challenging tasks in **Natural Language Processing (NLP)**. It is defined as the process of distilling the most important information from a source (or sources) to produce an abridged version for a particular user (or users) and task (or tasks) [139]. With the explosion in the amount of textual information available online since the advent of the Internet, summarization research has garnered significant attention.

This work was supported in part by the National Science Foundation under grants 2406251, 2201428, 2232862, 2004014, 2138259, IIS-2106972, III-2106758, and POSE-2346158. We thank NSF for enabling this research.

Authors' Contact Information: Haopeng Zhang, Department of Information and Computer Sciences, University of Hawaii at Manoa, Honolulu, United States; e-mail: workhaopeng@gmail.com; Philip S. Yu, University of Illinois at Chicago, Chicago, United States; e-mail: psyu@uic.edu; Jiawei Zhang, IFM Lab, Department of Computer Science, University of California, Davis, Davis, United States; e-mail: jiwzhang@ucdavis.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 0360-0300/2025/06-ART277

<https://doi.org/10.1145/3731445>




Fig. 1. The evolution of the four major paradigms in text summarization research.

Early efforts in building **automatic text summarization (ATS)** systems date back to the 1950s [134]. Subsequently, unsupervised feature-based systems emerged with advances in statistical machine learning during the 1990s and early 2000s [20, 53, 143]. In the 2010s, the focus of summarization research shifted toward training deep learning frameworks in a supervised manner, leveraging the availability of large-scale training data [28, 150, 153, 181]. More recently, the advent of self-supervised **pre-trained language models (PLMs)** such as **Bidirectional Encoder Representations from Transformers (BERT)** [43] and T5 [170] has significantly enhanced summarization performance through the ‘pre-train, then fine-tune’ pipeline [125, 131, 271]. This progression culminates in the current era dominated by **large language models (LLMs)** [13, 157]. Looking back at the development history of summarization approaches, we can generally categorize it into four stages based on the underlying paradigms: the statistical stage, the deep learning stage, the PLM fine-tuning stage, and the current LLM stage, as illustrated in Figure 1.

Recently, the emergence of LLMs has revolutionized both academic NLP research and industrial products due to their remarkable ability to understand, analyze, and generate texts with vast amounts of pre-trained knowledge. By leveraging extensive corpora of text data, LLMs can capture complex linguistic patterns, semantic relationships, and contextual cues, enabling them to produce high-quality summaries that rival those crafted by humans [261]. There is no doubt that LLMs have propelled the field of summarization into a new era [157, 201, 203].

At the same time, many NLP researchers are experiencing an existential crisis triggered by the astonishing success of LLM systems. Text summarization is definitely one of the most affected fields in which researchers argue that summarization is (almost) dead [166]. After such a disruptive change to our understanding of the summarization field, what is left to do? The ever-evolving landscape of LLMs, characterized by the continual development of larger, more powerful models, poses both opportunities and challenges for researchers and practitioners in the field of summarization.

This survey article aims to provide a comprehensive overview of the state-of-the-art summarization research endeavors in the new era of LLMs. We first categorize existing approaches and discuss the problem formulation, evaluation metrics, and commonly used datasets. Next, we systematically analyze representative text summarization methods prior to the LLM era, encompassing traditional statistical methods, deep learning approaches, and PLM fine-tuning techniques. Furthermore, we synthesize insights from recent summarization literature during the LLM era, discuss research trends and open challenges, and propose promising research directions in summarization. This survey seeks to foster a deeper understanding of the advancements, challenges, and future prospects in leveraging LLMs for text summarization, ultimately contributing to the ongoing development and refinement of NLP research.

1.1 Major Differences

Several surveys have been conducted to examine various aspects of summarization. Nevertheless, existing surveys mostly focus on investigating traditional statistical approaches and deep learning-based summarization approaches [15, 46, 52]. The scarcity of comprehensive up-to-date survey studies and the lack of consensus continue to inhibit progress. With the significant paradigm shift

initiated by PLMs and, more recently, LLMs, there is still a dearth of thorough investigations that comprehensively encompass the continuous progress in the field of summarization in this new era.

For example, early survey papers [52, 146] provide a comprehensive survey of statistical and deep learning-based ATS models, discussing their detailed taxonomy and applications. Following this, Cajueiro et al. [15] and Wibawa et al. [218] provide literature reviews of the methods, data, evaluation, and coding of ATS up to the advent of PLMs. In contrast to these surveys, our article presents a more structured categorization of ATS methods based on paradigm shifts, offering a clearer taxonomy that facilitates better organization, traceability, and understanding of the evolution of summarization techniques.

Researchers have also produced surveys with a more specific focus. Survey paper [94] focuses on the datasets, methods, and evaluation metrics of long documents such as academic articles and business reports. Another work [175] focuses on **graph neural network-based (GNN-based)** approaches to the task of ATS. Cao [16] conducts a survey on neural abstractive summarization methods and explores the factual consistency of the abstractive summarization systems.

Considering the rapid advancement of summarization approaches and the disruptive change introduced by LLMs, we believe it is imperative to review the details of representative methods both before and during the LLM era, analyze the uniqueness of each method, and discuss open challenges and prospective directions to facilitate the further advancement of the field.

1.2 Organization

This survey is organized as follows: Section 2 provides the background of text summarization, including categorizations of approaches, problem formulation, evaluation metrics, and common datasets. Section 3 reviews prominent summarization methods before the era of LLMs, encompassing statistical approaches, deep learning-based approaches, and PLM fine-tuning approaches. Section 4 explores recent advancements in benchmarking LLMs for summarization (Section 4.1), developing LLM-based summarization systems (Section 4.2), and evaluating summaries with LLMs (Section 4.3). Finally, Section 5 discusses open problems and future research directions in summarization.

2 Background

2.1 Categorization

Text summarization is the task of creating a short, accurate, and fluent summary of a longer text document. It finds applications across various domains such as news aggregation, document summarization, social media analysis, and more. Its primary goal is to help users quickly grasp the main points of a document or a piece of content without having to read through the entire text. However, text summarization faces challenges such as maintaining coherence and preserving important details, dealing with various types of content, and ensuring the summary is both accurate and concise. As depicted in Figure 2, text summarization approaches can be classified into various groups depending on the format of the input document, the style of the summary output, and the underlying paradigm.

2.1.1 Input: Single-Document vs. Multi-Document vs. Query-Focused. Common text summarization approaches can be categorized based on the format of the input source document into **single-document summarization (SDS)**, **multi-document summarization (MDS)**, and **query-focused summarization (QFS)**. SDS summarizes a single article [131], while MDS takes a cluster of documents, typically on the same topic, as input [54]. QFS, on the other hand, aims to generate a summary that specifically addresses an input query, such as a topic, keywords, or an entity [6, 272]. Additionally, some summarization tasks involve multilingual inputs (e.g., translating from Chinese to English) [72] and multimodal inputs (e.g., combining text and images) [83, 117].




Fig. 2. Categorization of summarization approaches based on input formats and output styles.

2.1.2 Output: Extractive vs. Abstractive vs. Hybrid. Text summarization approaches can also be classified into extractive, abstractive, and hybrid methods based on how the output summary is generated. Extractive methods create summaries by extracting sentences from the original documents [131]. Abstractive methods generate the summary word by word with novel content [107]. Hybrid methods combine both extraction and abstraction techniques [49].

Specifically, *extractive summarization* systems aim to identify and select significant text spans and sentences from the original document to form the summary. These summaries are inherently faithful, fluent, and accurate, but they may suffer from issues such as redundancy and incoherence. Statistical and deep learning-based summarization systems [53, 143, 150] are mostly extractive and typically frame extractive summarization as a sequence labeling and ranking problem to identify the most salient sentences.

Abstractive summarization systems, in contrast, generate summaries from scratch, similar to how humans write summaries. These systems produce more flexible and concise summaries but may encounter problems like hallucination and unfaithfulness. Abstractive summarization is often formulated as a sequence-to-sequence (seq2seq) problem [195] and utilizes encoder-decoder frameworks [107, 258] or autoregressive language models as in [13].

Hybrid summarization approaches attempt to combine the strengths of both extractive and abstractive methods to achieve more balanced results [49, 254]. These systems typically extract key information from the source document first and then use this extracted information to guide the abstractive summary generation process.

2.1.3 Backbone Paradigm: Statistical vs. Deep Learning vs. PLM Fine-Tuning vs. LLM. As previously mentioned, the development of text summarization can be categorized into four major stages based on the underlying paradigms: the statistical stage, the deep learning stage, the PLM fine-tuning stage, and the current LLM stage.

Statistical Stage features unsupervised methods for summarization, including heuristic-based approaches [20], optimization-based approaches [116, 142], and graph-based approaches [53, 143]. Hand-crafted features and frequency-based features like the **term frequency-inverse document frequency (TF-IDF)** are commonly used [51].

Deep Learning Stage approaches rely on domain-specific training data in the form of source document-summary pairs to train neural networks in a supervised manner [28, 150, 153, 181]. Word embedding techniques [144, 161] are also commonly used. During this stage, researchers introduced various training corpora to advance summarization development [73, 152, 176].

PLM Fine-tuning Stage: This stage relies on advances in large-scale, self-supervised PLMs [107, 123]. The introduction of BERT [43] marks the beginning of this stage, resulting in significant performance improvements. Pre-training with large amounts of text data equips PLMs with language patterns and parametric knowledge, which are important for document comprehension and text generation, thus benefiting downstream tasks. This stage features a ‘pre-train, then fine-tune’ pipeline that fine-tunes PLMs on task-specific data to further improve performance on downstream tasks [131].

LLM Stage: Most recently, advances in LLMs have reshaped summarization research. The strong understanding and instruction-following capabilities of these models have enabled the development of zero-shot and few-shot summarization systems, bringing new opportunities in this emerging era [13, 216]. The beginning of this stage is marked by the introduction of **Generative Pre-trained Transformer 3 (GPT-3)** [13] by OpenAI, with a model size of 175 billion parameters and strong few-shot capabilities.

2.2 Evaluation Metrics

Early research focuses on evaluating the quality of summaries using human judges [41, 139, 156]. Since manually evaluating generated summaries is both costly and impractical, researchers have developed automatic metrics instead. Generally, the overall quality of a summary is evaluated along four dimensions: coherence, consistency, fluency, and relevance [273]. Recently, Fabbri et al. [55] conduct a comprehensive meta-evaluation of common automatic evaluation metrics, incorporating crowd-sourced human annotations.

2.2.1 Similarity-Based Summary Evaluation. Existing summary evaluation metrics primarily rely on the similarity between the generated summary and a reference gold summary (human-written) as the main criterion.

ROUGE F-scores [115] have long been the standard way to evaluate summarization model performance. They measure the n-gram lexical overlap between the reference and candidate summaries. Specifically, ROUGE-1 and ROUGE-2 scores refer to unigram and bigram overlap, respectively, indicating summary informativeness. The ROUGE-L score, which refers to the longest common sequence, indicates summary fluency. However, since ROUGE is based on exact n-gram matches, it ignores overlaps between synonymous phrases and penalizes models that generate novel wordings and phrases.

Researchers have also explored contextualized embedding-based similarity metrics to evaluate summary quality, such as BERTScore [260], MoverScore [268], and Sentence Mover's Similarity [38]. These metrics heavily depend on the pre-trained encoder to vectorize the candidate and ground truth summaries, potentially introducing inherent biases and suffering from low interpretability.

On the other hand, Belz and Gatt [9] argue that this similarity reflects 'human-likeness' and may be unrelated to the final performance on generation tasks. BartScore [248] then conceptualizes the evaluation of generated text as a text generation problem. It argues that models trained to convert the generated text to/from a reference output or the source text will achieve higher scores when the generated text is better. Meanwhile, researchers are also exploring reference-free automatic evaluation metrics such as SummaQA [179], BLANC [205], and SUPERT [62].

Despite these advances, evaluation methodologies still lack comprehensive assessments across different domains. For example, news summarization relies on general coherence and informativeness, whereas clinical and scientific summarization requires factual precision and citation accuracy. A recent study also proposed a new benchmark to study the domain generalization capability of summarization systems [247]. Addressing these gaps, recent studies advocate for multi-dimensional evaluation frameworks that combine lexical, semantic, and factuality assessments tailored to the specific needs of different application areas.

2.2.2 Factual Consistency. The factual consistency of generated summaries is crucial for their real-world applications [98, 141, 249]. Researchers have delved into metrics to automatically evaluate summary faithfulness based on text entailment or **question answering (QA)**.

The text entailment-based approach evaluates factual inconsistency by verifying the summary against the original document. FactCC [98] is a weakly supervised BERT-based model metric that uses rule-based transformations applied to source document sentences. DAE [68] decomposes

entailment at the level of dependency arcs, examining the semantic relationships within the generated output and input. SummaC [100] segments documents into sentence units and aggregates scores between pairs of sentences.

QA-based metrics, on the other hand, employ a question-generation model to generate questions from a given summary and evaluate factual consistency by the extent to which the summary provides sufficient information to answer these questions. FEQA [50] and QAGS [210] measure summary factual consistency by determining how well a summary answers questions posed on its corresponding source document. Questeval [180] unifies precision and recall-based QA metrics, measuring summary factual consistency by assessing how many generated questions can be answered by the summary.

2.2.3 Coherence and Redundancy. SNaC [67] is a narrative coherence evaluation framework rooted in fine-grained annotations, specifically designed for long summaries. Peyrard et al. [162] propose measuring the redundancy extent of a summary using its unique n-gram ratio, which is the percentage of unique n-grams. Xiao and Carenini [227] introduce another metric to measure summary redundancy as the inverse of a diversity metric with length normalization, where diversity is defined as the entropy of unigrams in the document. However, coherence and redundancy are particularly challenging in LLM-generated summaries, as these models tend to introduce subtle inconsistencies while maintaining high fluency.

Comparing all the evaluation metrics, ROUGE remains a widely used metric due to its simplicity, interpretability, and alignment with traditional extractive summarization, but it struggles with capturing paraphrases, semantic equivalence, and factual accuracy, making it less effective for evaluating abstractive and LLM-based summarization. Embedding-based metrics such as BERTScore and MoverScore address some of these limitations by leveraging contextual representations, yet they introduce potential biases from pre-trained models and often lack transparency. Entailment-based metrics like FactCC and SummaC provide a structured way to assess factual consistency, but they are sensitive to noise in training data and may misclassify factually correct yet novel information as incorrect. QA-based methods, including FEQA and Questeval, offer a more interactive evaluation by verifying whether summaries contain sufficient information to answer generated questions; however, their effectiveness depends on the quality of the QA models, which can introduce propagation errors. Given these tradeoffs, recent research advocates for hybrid evaluation frameworks that integrate lexical overlap, semantic similarity, factual consistency, and domain-specific criteria to provide a more holistic assessment of summarization quality, particularly in the era of LLMs.

2.3 Summarization Datasets

The availability of public datasets for text summarization has been a major driving force behind the rapid advancement in recent years. These datasets vary across multiple dimensions, including domain, formats, size, and the number of gold summaries they provide. Here, we summarize the characteristics of these datasets and provide an overview in Table 1, which includes details such as the standard train-validation-test split, language, domain, summary format, and URL links. Our primary focus is on datasets containing summaries of texts written in English.

2.3.1 Short News Summarization Dataset. CNN/DM [73] is the most widely-used dataset for summarization. It contains news articles and associated highlights as summaries written by journalists from CNN and the DailyMail. The original dataset is created for machine comprehension and adapted for summarization by [149]. XSum [152] is a dataset consisting of one-sentence summaries derived from online articles published by BBC. However, it has been noted that some summaries in the XSum dataset suffer from low faithfulness, and contain information that cannot

Table 1. Overview of Common Text Summarization Datasets

Dataset	Size	Language	Domain	Format	Source	Link
CNN/DM	287, 084/13, 367/11, 489	English	News	SDS	[73]	url
XSum	203, 028/11, 273/11, 332	English	News	SDS	[152]	url
NYT	137, 778/17, 222/7, 223	English	News	SDS	[176]	url
NEWSROOM	995, 041/108, 837/108, 862	English	News	SDS	[69]	url
Gigaword	3, 803, 957/189, 651/1, 951	English	News	SDS	[174]	url
CCSUM	1, 349, 911/10, 000/10, 000	English	News	SDS	[85]	url
WikiHow	168, 126/6, 000/6, 000	English	Knowledge Base	SDS	[96]	url
Reddit	41, 675/645/645	English	Social Media	SDS	[90]	url
SAMSum	14, 732/818/819	English	Dialogue	SDS	[65]	url
MediaSum	463, 596	English	Dialogue	SDS	[278]	url
AESLC	14, 436/1, 960/1, 906	English	Email	SDS	[259]	url
ArXiv	201, 427/6, 431/6, 436	English	Scientific Paper	SDS	[39]	url
PubMed	112, 291/6, 402/6, 449	English	Scientific Paper	SDS	[39]	url
BIGPATENT	1, 207, 222/67, 068/67, 072	English	Patent	SDS	[186]	url
BillSum	18, 949/1, 237/3, 269	English	Bill	SDS	[95]	url
FINDSum	42, 250	English	Report	SDS	[122]	url
DUC 05/06/07	50 × 32/25/10	English	News	MDS	[159]	url
MultiNews	44, 972/5, 622/5, 622	English	News	MDS	[54]	url
WikiSum	1.5m/38k/38k	English	Wikipedia	MDS	[121]	url
WCEP	8, 158/1, 020/1, 022	English	Wikipedia	MDS	[64]	url
Multi-XScience	30, 369/5, 066/5, 093	English	Scientific Paper	MDS	[133]	url
Yelp	1, 038, 184/129, 856/129, 840	English	Review	MDS	[36]	url
QMSum	1, 257/272/279	English	Meeting	QFS	[272]	url
NewTS	4800/-/1200	English	News	QFS	[6]	url
TD-QFS	3, 400	English	Medical	QFS	[7]	url
XL-Sum	1, 005, 292	Multiple	News	SDS	[72]	url

be directly inferred from the source document [141]. **NYT** [176] contains news articles with abstractive summaries from the New York Times. Input documents are commonly truncated to 768 tokens following [263]. **NEWSROOM** [69] is a large dataset containing article-summary pairs written by authors and editors in the newsrooms of 38 major publications between 1998 and 2017. **Gigaword** [174] contains around 4 million headline-article pairs extracted from news articles (seven publishers) from the Gigaword corpus. **CCSUM** [85] is a more recent news summary dataset that contains 1.3 million high-quality recent training news articles, filtered from CommonCrawl News articles. **XL-Sum** [72] is a comprehensive and diverse dataset comprising 1 million professionally annotated article-summary pairs from BBC. The dataset covers 44 languages and is highly abstractive, concise, and of high quality.

2.3.2 Other Domains Short Summarization Dataset. **WikiHow** [96] is a diverse dataset extracted from online knowledge base. Each article consists of a number of steps, and the gold-standard summary for each document is the concatenation of bold statements. **Reddit** [90] contains 120K posts of informal stories from the TIFU online discussion forum subreddit. The summaries are highly abstractive and informal. **SAMSum** [65] consists of approximately 16,000 single-document, text message-like dialogues. These documents are created by linguists fluent in English and cover a broad spectrum of formality levels and discussion topics. **AESLC** [259] is an email domain dataset for the task of automatically generating an email subject line from the email body.

2.3.3 Long Document Summarization Dataset. **PubMed** and **arXiv** [39] are two widely-used long document datasets of scientific publications from arXiv.org and PubMed. The task is to generate

the abstract of the paper from the body content. **BIGPATENT** [186] consists of 1.3 million US patents along with human summaries under nine patent classification categories. **BillSum** [95] contains 23k US Congressional and California state bills and human-written reference summaries from the 103rd–115th (1993–2018) sessions of Congress. **FINDSum** [122] is a new large-scale dataset for long text and multi-table summarization built on 21,125 annual reports from 3,794 companies.

2.3.4 Multi-Document Summarization Dataset. **DUC** [159] presents an overview of the datasets provided by the **Document Understanding Conferences (DUC)**. The DUC datasets from 2005 to 2007 include MDS datasets in the news domain, with 10–30 documents and 3–4 human-written summaries per cluster. **MultiNews** [54] is a large-scale multi-document news summarization dataset from the newser.com website with human-written summaries. Each summary is professionally written by editors, and the dataset covers diverse news sources. **WikiSum** [121] is a MDS dataset where each summary is a Wikipedia article and the source documents are either citations in the reference section or the web search results of section titles. **WCEP** [64] is built based on news events from **Wikipedia Current Events Portal (WCEP)** and similar articles in the CommonCrawl News dataset. It features high alignment with several real-world industrial use cases and has 235 articles per cluster on average. **Multi-XScience** [133] is a MDS dataset created from scientific articles. The summaries are paragraphs of the related work section, while source documents include the abstracts of the query and referred papers. **Yelp** [36] is a customer review dataset created from the Yelp Dataset Challenge. The task is to summarize multiple reviews of a business.

2.3.5 Query-Focused Summarization Dataset. **QMSum** [272] is a meeting summarization dataset that consists of 1,808 query-summary pairs over 232 meetings in multiple domains. The dataset also contains the main topics of each meeting and the ranges of relevant text spans for the annotated topics and each query. **NewTS** [6] is an aspect-focused summarization dataset derived from the CNN/DM dataset and contains two summaries focusing on different topics for the same news. **TD-QFS** [7] is centered on medical texts and comprises short keyword queries across 4 clusters, each containing 185 documents. The clusters in TD-QFS exhibit lower topical concentration, featuring larger quantities of query-irrelevant information.

3 Summarization Methods Prior to LLMs

Before delving into the new opportunities enabled by recent advancements in LLMs, this section reviews the significant progress previously achieved in summarization by investigating representative summarization methods and algorithms, as outlined in the overview taxonomy shown in Figure 3. This comprehensive discussion encompasses the three major summarization paradigms prior to LLMs: statistical methods (Section 3.1), deep learning-based approaches (Section 3.2), and PLM-based fine-tuning approaches (Section 3.3). The details, including method category, backbone, and datasets of these representative methods, are summarized in Table 2.

3.1 Statistical Summarization Methods

In the early stages of text summarization systems, research primarily focuses on extractive summarization and relies on traditional statistical approaches to automatically create summaries. These systems mostly use frequency-based features such as TF-IDF and hand-crafted features [51] to model text data. Representative methods from this stage include heuristic-based methods, optimization-based methods, and graph-based ranking methods.

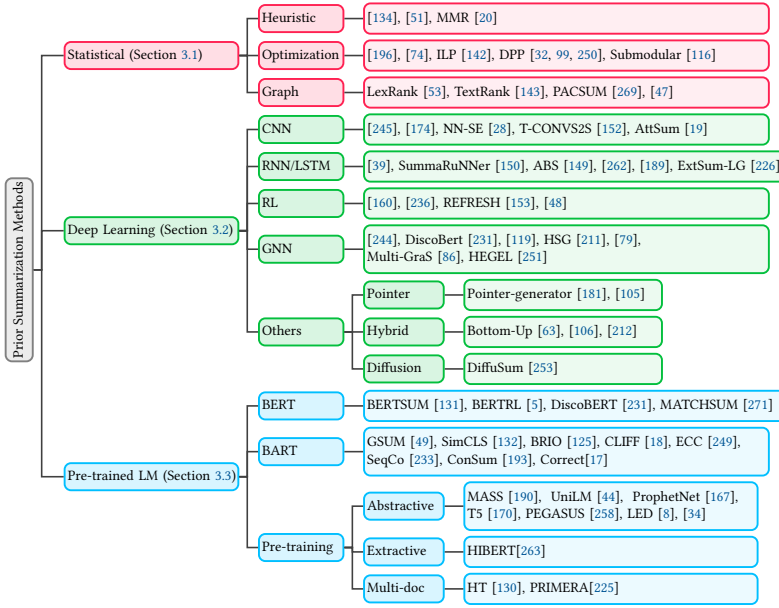


Fig. 3. Taxonomy of representative summarization methods prior to LLMs.

3.1.1 Heuristic-Based Methods. One of the earliest summarization approaches is by Luhn [134]. This method hypothesizes that the importance of sentences can be measured by the frequency of specific content words (keywords). It scores sentences based on importance and then extracts the high-scoring ones to generate literature abstracts.

Edmundson [51] later proposes modeling sentences based on a combination of handcrafted features, including word frequency within the article, sentence position, the number of words in the article title or section headings, and the frequency of cue words (keywords). A simple linear summation is then applied to score and rank the sentences.

Maximal Marginal Relevance (MMR) [20] strives to reduce redundancy while maintaining query relevance in selecting sentences for extractive summarization. It employs a greedy approach to combine sentence relevance with information novelty.

3.1.2 Optimization-Based Methods. The sentence selection problem in extractive summarization has been formulated as a constrained optimization problem to obtain a globally optimal set of sentences. Summary sentences are selected to maximize the coverage of important source content while minimizing redundancy and adhering to a length constraint. Text summarization has been formalized as various optimization problems, such as the maximum coverage problem [196] and the tree knapsack problem [74], and addressed with optimization methods including **integer linear programming (ILP)** [142], **determinantal point process (DPP)** [32, 99], and submodular functions [116].

3.1.3 Graph-Based Methods. LexRank [53] and TextRank [143] are two seminal works that introduce the concept of representing documents as graphs and formulating extractive summarization as identifying the most central nodes in a graph, inspired by the PageRank algorithm [12]. The underlying premise is that the centrality of a node in the graph can serve as an approximation of the significance of its corresponding sentence within the document.

More recently, PACSUM [269] revisited the TextRank algorithm and proposed constructing graphs with directed edges. They argue that the relative position of nodes in a document influences

Table 2. Overview of Representative Summarization Methods before LLMs

Paradigm	Category	Backbone	Year	Method	Datasets
Statistical	UnSup. Ext.	Heuristic	1998	MMR [20]	TIPSTER topic
	UnSup. Ext.	Graph	2004	TextRank [143]	Inspec
	UnSup. Ext.	Graph	2004	LexRank [53]	DUC 03/04
	UnSup. Ext.	Optimization	2007	ILP [142]	DUC 02/05
	UnSup. Ext.	Optimization	2011	Submodular [116]	DUC 03-07
	UnSup. Ext.	Optimization	2012	DPP [99]	DUC 03/04
	UnSup. Ext.	Graph	2019	PACSUM [269]	CNNNDM, NTY, TTNEWS
DL	Sup. Ext.	CNN,LSTM	2016	NN-SE [28]	CNNNDM, DUC2002
	Sup. Ext.	RNN	2017	SummaRuNNer [150]	CNNNDM, DUC2002
	Sup. Abs.	Pointer network	2017	PTGEN [181]	CNNNDM
	Sup. Ext.	GNN	2017	GCN [244]	DUC
	Sup. Abs.	CNN	2018	T-CONVS2S [152]	CNNNDM, XSUM
	Sup. Hybrid	LSTM	2018	BottomUp [63]	CNNNDM, NYT
	Sup. Ext.	GRU	2018	NEUSUM [275]	CNNNDM
	RL Ext.	RL	2018	REFRESH [153]	CNNNDM
	RL Ext.	RL	2018	Intra-attention [160]	CNNNDM, NYT
	Sup. Ext.	LSTM	2019	ExtSum-LG [226]	PubMed, ArXiv
	Sup. Ext.	GNN	2020	HSG [211]	CNNNDM, NYT, MultiNews
	Sup. Ext.	GNN	2022	HEGEL [251]	arXiv, PubMed
	Sup. Ext.	Diffusion	2023	DiffuSum [253]	CNNNDM, XSUM, PubMed
PLM	Sup. Abs.	Transformer	2019	HT [130]	WikiSum
	Sup. Ext.	Transformer	2019	HIBERT [263]	CNNNDM, NYT
	Sup. Abs.	Transformer	2019	MASS [190]	CNNNDM, XSUM, Gigaword
	Sup. Ext.	BERT	2019	BERTSUMEXT [131]	CNNNDM, NTY, XSUM
	Sup. Abs.	BERT	2019	BERTSUMABS [131]	CNNNDM, NTY, XSUM
	Sup. Abs.	Transformer	2019	UniLM [44]	CNNNDM, Gigaword
	Sup. Abs.	Transformer	2019	BART [107]	CNNNDM, XSUM
	Sup. Abs.	Transformer	2019	T5 [170]	CNNNDM
	Sup. Ext.	Graph, BERT	2020	DiscoBERT [231]	CNNNDM, NYT
	Sup. Ext.	BERT	2020	MATCHSUM [271]	CNNNDM, XSUM, etc.
	Sup. Abs.	Transformer	2020	PEGASUS [258]	CNNNDM, XSUM, etc.
	Sup. Abs.	Transformer	2020	ProphetNet [167]	CNNNDM, Gigaword
	Sup. Abs.	Transformer	2020	LED [8]	arXiv
	Sup. Abs.	BART	2021	SimCLS [132]	CNNNDM, XSUM
	Sup. Hybrid	BERT, BART	2021	GSUM [49]	CNNNDM, XSUM, etc.
	Sup. Abs.	BART	2021	CLIFF [18]	CNN/DM, XSUM
	Sup. Abs.	BART	2022	BRIO [125]	CNNNDM, XSUM, NYT
	Sup. Abs.	LED	2022	PRIMERA [225]	WikiSum, arXiv, etc.

Unsup. stands for unsupervised, Sup. stands for supervised, Ext. stands for extractive, Abs. stands for abstractive.

the contribution of any two nodes to their centrality. Additionally, PACSUM integrates neural representations to improve sentence similarity computation.

3.2 Deep Learning Summarization Methods

Advances in deep neural networks and the availability of large-scale text corpora have significantly accelerated progress in various NLP tasks. The development of word embedding techniques, such as Word2Vec [144] and GloVe [161], has enabled deep learning models to represent text data as sequences of tokens, automatically learning continuous sentence and document representations from data and eliminating the need for human-engineered features. Common deep learning frameworks used in NLP include **convolutional neural networks (CNNs)**, **recurrent neural networks**

(RNNs), **long short-term memory (LSTM)** networks [75], **reinforcement learning (RL)**, and GNNs. Summarization research in the deep learning stage still predominantly focuses on extractive summarization in a supervised manner. However, abstractive summarization methods are beginning to emerge, leveraging frameworks such as sequence-to-sequence (seq2seq) models [195], pointer networks [208], and attention mechanisms.

3.2.1 CNN-Based Methods. After the successful application of CNNs for learning continuous document representations in text classification tasks [91], Yin and Pei [245] propose a CNN-based neural network language model. This model projects sentences into dense distributed representations and employs a diversified selection process to extract fewer redundant sentences for summaries. Rush et al. [174] apply CNN for abstractive summarization with the sequence-to-sequence model structure. It uses CNN to encode the source and a context-sensitive attentional feed-forward neural network to generate the summary.

NN-SE [28] introduces a hierarchical document encoder that combines a convolutional sentence encoder for learning sentence representations with a recurrent document encoder for learning document representations. Additionally, NN-SE incorporates an attention-based neural sentence extractor to select sentences or words for summarization.

T-CONVS2S [152] proposes a topic-aware convolutional sequence-to-sequence learning approach for the extreme summarization task. The distributions obtained from **Latent Dirichlet Allocation (LDA)** [11] are used as additional input. The convolutional encoder associates each word with a topic vector, while the convolutional decoder conditions each word prediction on a document topic vector.

For QFS, AttSum [19] proposes a method that jointly ranks query relevance and sentence saliency. It automatically learns distributed representations for sentences and queries within the latent space using CNNs.

3.2.2 RNN/LSTM-Based Methods. SummaRuNNer [150] is a representative RNN-based extractive summarization model. It employs a two-layer bidirectional RNN, with the bottom layer operating at the word level within each sentence and the top layer running over sentences. An accumulated document representation is generated through a weighted sum of sentence representations, and a logistic classifier is used to make predictions.

ABS [149] first applies the attentional encoder-decoder RNN for abstractive summarization, proposing to model rare or unseen words with a switching generator-pointer mechanism and to capture the document structure with hierarchical attention.

For long document summarization, Cohan et al. [39] propose a pioneering model for abstractive summarization of long documents. Their approach employs a hierarchical encoder to capture the discourse structure (scientific paper sections) of a document and an attentive discourse-aware decoder to produce the summary. Building on this, ExtSum-LG [226] proposes using the global context of the entire document along with the local context. It applies the LSTM-Minus method to capture the local context of each sentence.

3.2.3 RL-Based Methods. RL is a popular alternative for training abstractive and extractive summarization systems, addressing the exposure bias and objective mismatch issues associated with the **negative log-likelihood (NLL)** training loss [59].

A seminal work in this domain is REFRESH [153], which applies RL to extractive summarization by conceptualizing it as a sentence ranking task. This study introduces an objective function that integrates the maximum-likelihood cross-entropy loss with rewards derived from policy gradient RL. Paulus et al. [160] further advance this direction by introducing a neural network model with a novel intra-attention mechanism that attends to the input and the continuously generated output

separately. They also propose a new training method that combines standard supervised word prediction with RL. Moreover, Yadav et al. [236] propose two novel rewards obtained from the downstream tasks of question-type identification and question-focus recognition to regularize the generation model.

3.2.4 GNN-Based Methods. Researchers have also explored the use of GNNs [92, 207, 257] for NLP due to their ability to better capture the structural information of a document, as well as their scalability and interpretability [175, 242, 252, 256, 267]. GNNs typically use the adjacency matrix and initial node representation matrix as inputs, performing information aggregation on the graph through either spatial or spectral convolution. Graph Attention Networks (GATs) [207] is one of the most popular GNNs that introduces attention mechanism to the spatial convolution to filter unimportant neighbors. Here α_{ij} denotes the normalized attention scores between node i and node j .

Yasunaga et al. [244] is one of the pioneering endeavors in applying GNN to extractive summarization tasks. Their approach involves modeling documents as approximate discourse graphs and leveraging **Graph Convolutional Networks (GCN)** [92] to encode sentence representations, which are subsequently used for predicting sentence salience. Similarly, DiscoBert [231] introduces a method that employs GCN on structural discourse graphs constructed from **Rhetorical Structure Theory (RST)** trees and coreference mentions. Moreover, HSG [211] proposes an approach that views documents as heterogeneous graphs, incorporating semantic nodes of various granularity levels (words, sentences, documents). The model constructs word-document heterogeneous graphs and employs word nodes as the intermediary between sentences, enabling a more nuanced representation of document semantics.

Researchers have also explored various graph variants to enhance document modeling. MultiGraS [86] introduces the utilization of multiplex graphs to capture multiple types of inter-sentential relationships, such as semantic similarity and natural connection, while also modeling intra-sentential relationships. Furthermore, HEGEL [251] proposes a novel approach by representing documents as hypergraphs to capture high-order cross-sentence relations instead of pairwise relations. This method integrates diverse types of sentence dependencies, including section structure, latent topics, and keyword coreference, thereby offering a comprehensive view of document semantics.

3.2.5 Others. In addition to the aforementioned models, here we introduce some representative individual models.

See et al. [181] introduce the pointer-generator network [208] to mitigate the issue of inaccurate reproduction of factual details by seq2seq models. This approach enables the model to both generate words from the vocabulary using a generator and selectively copy content from the source using a pointer. Lebanoff et al. [105] further adapt the model to the multi-document summarization setting with MMR.

Bottom-Up [63] introduces a data-efficient content selector to identify potential output words in a source document. This selector is integrated as a bottom-up attention mechanism within the pointer-generator network, representing a hybrid approach that combines extractive and abstractive methods. Later, Lebanoff et al. [106] propose selecting and fusing sentences into summaries.

DiffuSum [253] proposes a method to extract summaries by directly generating desired summary sentence representations using diffusion models and extracting sentences based on matching these representations. The approach optimizes a sentence encoder with a matching loss for aligning sentence representations and a multi-class contrastive loss for promoting representation diversity.

3.3 Pre-Trained Language Model Summarization Methods

With the advent of Transformer-based architectures [206] and the powerful large-scale self-supervised pre-training, summarization systems have experienced a substantial leap in performance




Fig. 4. An illustration of the existing prevalent pre-training frameworks (taken from [107]). Encoder-only models mask random tokens during pre-training and encode documents bidirectionally. Decoder-only models predict tokens autoregressive conditioned on leftward context. Encoder-decoder models encode corrupted documents with a bidirectional model and then calculate generation likelihood with an autoregressive decoder.

over the past few years. The self-attention mechanism within the Transformer architecture facilitates parallel computation and enhances learning efficiency, allowing for effective training on large-scale unlabeled corpora. PLMs can learn universal language representations through this unsupervised training, offering superior model initialization for downstream tasks and eliminating the need to learn each new task from scratch. In this stage, summarization systems typically adapt PLM checkpoints, which have been trained in a self-supervised manner on large-scale corpora, and subsequently fine-tune these PLMs on domain- or task-specific training datasets.

As illustrated in Figure 4, three prevalent pre-training architectures are commonly used: encoder-only, encoder-decoder, and decoder-only. All these architectures employ the same self-attention layers as in vanilla Transformers to encode word tokens. The BERT [43] marks the beginning of this phase, being the first widely adopted PLM that can be used for both extractive and abstractive summarization with its encoder-only architecture [131]. Summarization research with PLMs has increasingly focused on abstractive summarization using encoder-decoder architecture models like T5 [170] and BART [107] as the backbone. OpenAI has introduced decoder-only models with generative pre-training [169]. Additionally, researchers have developed summarization-specific PLMs like PEGASUS [258], long-context PLMs like LED [8], and MDS PLMs like PRIMERA [225].

In this discussion, we use BERT as a representative for encoder-only PLMs and BART for encoder-decoder models to explore their applications in text summarization in Sections 3.3.1 and 3.3.2. We also examine representative efforts in pre-training from scratch in Section 3.3.3.

3.3.1 BERT-Based Methods. Encoder-only PLMs like BERT are predominantly used to convert natural language into continuous contextualized representations (embeddings), which can be further manipulated for downstream tasks. BERT is self-supervisedly pre-trained using the **next sentence prediction (NSP)** and **masked language modeling (MLM)** objectives on a large-scale unlabeled corpus. This encoder-only architecture makes BERT well-suited for various predictive modeling tasks, such as text classification and extractive summarization. Other popular encoder-only PLMs include RoBERTa [123], DistilBERT [177], ALBERT [102], and Sentence-BERT [173].

PreSum [131] stands as a seminal paper and the most popular baseline model for PLM-based summarization. It showcases how BERT could be effectively applied to text summarization and proposes a comprehensive framework for both extractive and abstractive summarization. They use several inter-sentence Transformer layers for extractive summarization and introduce a novel fine-tuning schedule that employs different optimizers for the encoder and the decoder, aiming to address the mismatch between these two components. Building upon PreSum, Bae et al. [5] further enhance it with RL techniques to achieve summary-level scoring. They propose bridging extractive and rewrite models using policy-based RL.

DiscoBERT [231] introduces an enhancement to BERT's document encoder by incorporating a graph encoder focused on structural discourse graphs. This augmentation aims to capture long-range contextual dependencies within the document.

MATCHSUM [271] introduces a novel approach to extractive summarization by framing the task as a semantic text-matching problem. Instead of extracting sentences individually, this method matches a source document with candidate summaries in a semantic space. The model uses a Siamese-BERT architecture, comprising two BERT models with tied weights, alongside a cosine-similarity layer during the inference phase.

3.3.2 BART-Based Methods. Encoder-decoder PLMs like BART are widely used for NLP tasks that involve understanding input sequences and generating output sequences due to their ability to capture the mapping between these sequences. BART is self-supervisedly pre-trained by denoising and reconstructing text that has been corrupted with an arbitrary noising function on a large-scale unlabeled corpus. The encoder-decoder architecture makes BART well-suited for various sequence-to-sequence generation tasks, such as machine translation and abstractive summarization. Other popular encoder-decoder PLMs include MASS [190], T5 [170], UL2 [200], and instruction-tuned versions such as Flan-T5 and Flan-UL2 [37].

Besides directly applying PLMs for abstractive summarization, Aghajanyan et al. [2] propose regularized fine-tuning rooted in trust-region theory to alleviate representational collapse, which is the degradation of generalizable representations of pre-trained models during the fine-tuning stage. GSUM [49] introduces a guided summarization framework that incorporates external guidance as additional input to improve the faithfulness and controllability of the generated summaries.

Additionally, several research efforts have also explored the use of contrastive learning in fine-tuning BART. SimCLS [132] introduces a two-stage framework where the BART model generates candidate summaries, and a RoBERTa model, fine-tuned as an evaluation model, scores and selects the best candidate summaries, leveraging contrastive learning. SeqCo [233] and ConSum [193] propose leveraging contrastive learning to improve BART fine-tuning, specifically to alleviate exposure bias. BRIO [125] introduces a novel training paradigm that assumes a non-deterministic distribution, assigning probability mass to different candidate summaries based on their quality.

On the other hand, research on adapting the BART model has begun to focus on fine-tuning techniques to improve the faithfulness of generated summaries. These techniques include error correction [17], contrastive training as implemented in CLIFF [18], post-processing methods [25], knowledge graph methods [140, 277], data filtering approaches [151], and further advancements in controlled generation as seen in ECC [249] and Spancopy [228].

3.3.3 Pre-Training Methods. Besides the efforts in fine-tuning PLMs on task-specific training data, researchers have also attempted to pre-train Transformer models with self-supervised objectives from scratch to better align with downstream tasks.

For extractive summarization, HIBERT [263] proposes a hierarchical BERT model for document encoding, pre-trained using document masking and sentence prediction with unlabeled data. It leverages the hierarchical structure of documents to improve the quality of the extracted summaries.

Early efforts in pre-training language models for abstractive summarization include masked seq2seq pre-training in MASS [190], unified pre-training in UniLM [44], and n-gram prediction pre-training in ProphetNet [167]. PEGASUS [258] is a popular PLM specifically designed for summarization. In PEGASUS, important sentences are removed or masked from an input document and are generated together as one output sequence from the remaining sentences, mimicking the process of creating an extractive summary.

Longformer-Encoder-Decoder (LED) [8] is a PLM specially designed for processing long documents. Longformer replaces the standard self-attention mechanism in Transformers with a combination of local windowed attention and task-motivated global attention. This modification allows the model to scale linearly with sequence length, enabling it to process documents consisting

of thousands of tokens more efficiently. Cho et al. [34] later extend LED with joint training with segmentation and Liu et al. [124] improve the sparse attention with heterogeneous graphs.

For MDS, HT [130] augments vanilla Transformers with hierarchical inter-paragraph attention and document graph representation. PRIMERA [225] proposes the first PLM specifically designed for multi-document representation, based on the LED architecture. It introduces a new gap sentence generation pre-training objective tailored for the multi-document setting, teaching the model to connect and aggregate information across documents using entity pyramid masking.

4 LLM-Based Summarization Research

Recently, the advent of LLMs like GPT-X has revolutionized the landscape of NLP and significantly influenced the direction of research in text summarization. Most current summarization studies now employ LLMs as backbone models, leveraging zero-shot or few-shot settings instead of the previous fine-tuning approach. Given LLMs' strong capabilities in generating coherent and contextually relevant text, LLMs have driven a shift toward abstractive summarization. This section provides a comprehensive overview of recent advancements in LLM-based summarization, highlighting key developments and emerging research trends.

The rise of LLMs has introduced a paradigm shift in text summarization, moving away from traditional supervised and unsupervised methods. Unlike traditional models that rely on task-specific labeled datasets for training, LLMs are pre-trained on vast corpora and can generate summaries through zero-shot and few-shot learning, reducing the dependency on manually curated summarization datasets. Furthermore, LLMs exhibit enhanced contextual understanding, leveraging broad world knowledge and dynamic information integration to produce more coherent and contextually relevant summaries. This shift enables LLM-based approaches to generalize across diverse domains, such as news, scientific literature, and clinical summarization, without extensive domain-specific fine-tuning. Additionally, techniques such as instruction tuning [158] and **reinforcement learning from human feedback (RLHF)** Christiano et al. [35] further refine LLM-generated summaries, enhancing their adaptability and alignment with user preferences.

Despite these advantages, LLM-based summarization also presents distinct challenges. The vast number of parameters in LLMs makes training and adapting them for domain-specific summarization computationally demanding. Consequently, researchers are investigating parameter-efficient tuning methods such as LoRA [77] to mitigate these challenges. Furthermore, domain-specific evaluation metrics are becoming increasingly important, particularly in fields like clinical and scientific summarization, where factual accuracy and citation consistency are critical. Addressing these challenges is essential for maximizing the effectiveness and reliability of LLM-based summarization systems.

An overview of the current research landscape in LLM-based summarization is presented in Figure 5. Existing research works can be broadly categorized into three distinct types based on the objective of the study: LLM benchmarking studies, summarization modeling studies, and summary evaluation studies:

- **Benchmarking studies:** These studies focus on evaluating the performance of LLMs across various summarization tasks. They aim to establish baselines and performance metrics that help in comparing different models and approaches. Benchmarking studies are crucial for understanding the strengths and limitations of LLMs in summarization and for guiding future research directions. They often involve extensive experimentation and analysis across diverse datasets and domains, providing a comprehensive performance landscape of existing models. Representative benchmarking studies are presented in Section 4.1 and summarized in Table 3.

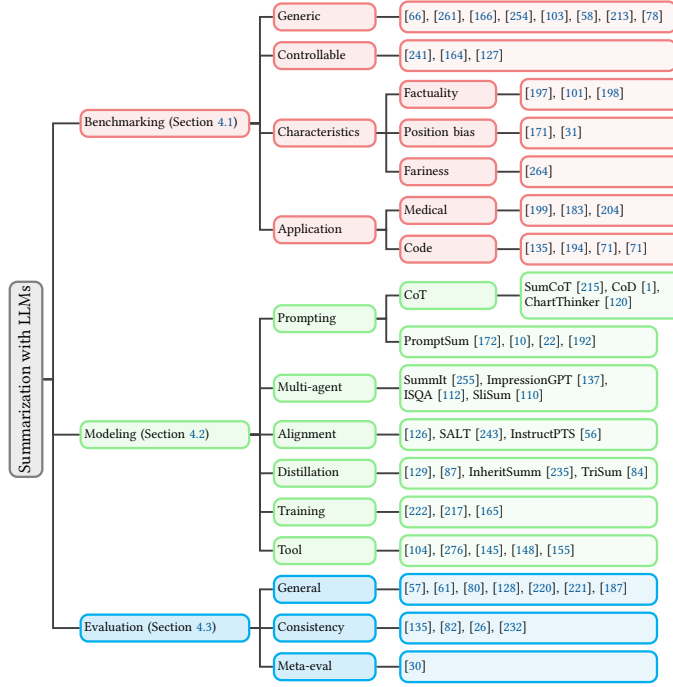


Fig. 5. Taxonomy of summarization research with LLMs.

- **Modeling Studies:** This category encompasses research dedicated to developing novel summarization algorithms and architectures utilizing LLMs. These studies explore various techniques to enhance the summarization process, such as prompt engineering, multi-agent systems, model alignment, and knowledge distillation. Modeling studies aim to push the boundaries of what LLMs can achieve in summarization, focusing on improving the quality, coherence, and informativeness of generated summaries. They also investigate how different architectural modifications and training strategies can optimize the performance of LLMs for specific summarization tasks. Representative modeling studies are presented in Section 4.2 and summarized in Table 4.
- **Evaluation Studies:** These studies are concerned with devising new metrics and methods for assessing the quality and effectiveness of generated summaries. Unlike traditional metrics like ROUGE, evaluation studies in the era of LLMs focus on developing more sophisticated and human-aligned metrics. Representative evaluation studies are presented in Section 4.3 and summarized in Table 5.

4.1 Benchmarking Studies

With the advent of powerful new LLMs as backbone models, many research efforts are now aimed at benchmarking their performance and understanding their capabilities in the context of summarization. Existing works have explored benchmarking LLMs for generic and controllable summarization, examining characteristics such as faithfulness and position bias, and investigating their application in specialized domains such as medical and code summarization.

4.1.1 Generic Summarization. The emergence of GPT-3 marks a pivotal moment in the LLM era, demonstrating significant generative capabilities that spur a paradigm shift in NLP research. One

Table 3. Overview of Benchmarking Studies on Summarization with LLMs

Source	Year	Task/Domain	LLM	Dataset
[66]	2022	News	GPT-3	CNN/DM, XSum, Newsroom
[261]	2023	News	GPT-3	CNN/DM, XSum
[166]	2023	News	GPT-4	CNN/DM, Multi-News, Mediasum
[254]	2023	Extractive	ChatGPT	CNN/DM, XSum, Reddit, PubMed
[103]	2023	Meeting	multiple	AMI, ICSI, QMSUM
[58]	2024	Meeting	Multiple	In-Domain, QMSUM
[213]	2023	Multilingual	Multiple	CLS
[78]	2023	Multi-doc	Multiple	DIVERSESUMM
[241]	2023	QFS	ChatGPT	CovidET, NEWTS, QMSum, SQuALITY
[164]	2023	QFS	ChatGPT	ELIFE
[127]	2023	Controllable	Multiple	INSTRUSUM
[197]	2023	Factuality	Multiple	CNN/DM, XSum
[101]	2023	Factuality	Multiple	AggreFact, DialSummEva
[198]	2023	Factuality, dialogue	Multiple	MediaSum, MeetingBank
[264]	2023	Fairness	Multiple	Claritin, US Election, Yelp, etc.
[171]	2023	Position bias	Multiple	MiddleSum
[31]	2024	Position bias	Multiple	CNN/DM, XSum, News, Reddit
[199]	2023	Medical	ChatGPT	Alzheimer, Kidney, Skin, etc.
[183]	2023	Medical	GPT-3	RCT
[204]	2023	Medical	Multiple	ProbSum, MeQSum, ACI-Bench, etc.
[135]	2023	Code	Multiple	BinSum
[194]	2023	Code	ChatGPT	CSN-Python
[71]	2024	Code	CodeT5, Llama, PaLM 2	CodeXGLUE

of the pioneering benchmarking studies [66] evaluates GPT-3 on text summarization, specifically within the classic benchmark domain of news summarization. This study compares GPT-3 against previously fine-tuned models and finds, through human evaluation, that annotators overwhelmingly preferred GPT-3's summaries, even when provided with only a task description. However, this preference is not reflected in automatic metrics, where GPT-3's summaries receive lower scores in both reference-based and reference-free evaluations.

Further large-scale human evaluation by Zhang et al. [261] echoes these findings, highlighting limitations due to low-quality reference summaries. These limitations lead to underestimations of human performance and reduced few-shot and fine-tuning effectiveness. Their study underscores that LLM-generated summaries are judged to be comparable to human-written ones and emphasizes the significance of instruction tuning in enhancing the zero-shot summarization capabilities of LLMs. Similarly, human evaluation [166] reaches similar conclusions, arguing that LLM-generated summaries are noted for better factual consistency and fewer instances of extrinsic hallucinations.

For extractive summarization, Zhang et al. [254] conduct a thorough evaluation of ChatGPT's extractive performance against traditional fine-tuning methods. Their findings reveal that while ChatGPT underperforms in terms of ROUGE scores compared to existing supervised systems, it achieves higher performance based on LLM-specific evaluation metrics. The study also explores the benefits of **in-context learning (ICL)** and proposes an extract-then-generate pipeline to improve the faithfulness of LLM-generated summaries. Here ICL is a method of prompt engineering that allows LLMs to learn new tasks without fine-tuning by using examples as part of a prompt [45].

For meeting summarization, Laskar et al. [103] further conduct an evaluation of LLMs on meeting summarization, discussing the cost-effectiveness and computational demands. This study provides a practical perspective on the tradeoffs involved in using LLMs in real-world applications, highlighting

Table 4. Overview of LLM-Based Summarization Methods

Source	Year	Task/Domain	LLM	Dataset
[172]	2023	Book	PEGASUS	CNN/DM, XSum, BillSum, SAMSum
[10]	2022	Opinion	GPT-3.5	SPACE, FewSum
[22]	2023	Book	GPT-4, Claude 2, etc,	-
[192]	2024	News	GPT-3.5/4	InstruSum
[215]	2023	News	GPT-3	CNN/DM, XSum
[1]	2023	News	GPT-4	CNN/DM
[120]	2024	Chart	LLaMa2	Chart-Sum-QA
[255]	2023	News	ChatGPT	CNN/DM, XSum, NEWTS
[112]	2024	Scientific	Llama2, Vicuna, Flan-T5	SciMRC, QASPER
[137]	2023	Radiology report	ChatGPT	MIMIC-CXR, OpenI
[110]	2024	multiple	Llama2, Claude-2, GPT-3.5	CNN/DM, XSum, arXiv, PubMed
[126]	2022	News	GPT-3.5, T5	DEFACTO
[243]	2023	Conversation	GPT-2	CC
[56]	2023	Title	Flan T5	-
[129]	2023	News	ChatGPT,GPT-4	CNN/DM, XSum
[235]	2023	Multiple	GPT-3.5, ZCode++	MultiNew, SAMSum, etc.
[87]	2023	Multiple	ChatGPT, GPT-3.5, T5	Gigaword
[84]	2023	News, clinical	GPT-3.5	CNN/DailyMail, XSum, ClinicalTrial
[222]	2024	News	Flan-T5, BART	CNN/DM, MultiNews, Gigaword
[165]	2024	Long doc	Vicuna, GPT-4	Multi-LexSum, eLife, BookSum
[104]	2023	QFS	GPT-4	Debatepedia
[276]	2023	Dialogue	ChatGPT	SAMSum, DIALOGSUM
[148]	2023	News	GPT-3.5	KEYEVENTS
[145]	2023	Dialogue	GPT-3.5	Tweetsumm

Table 5. Overview of LLM-Based Summary Evaluation

Source	Year	Aspect	LLM
[57]	2023	Coverage, Factuality, Coherence, Informativeness, etc.	GPT-3, OPT, FLAN-T5, GPT-2
[61]	2023	Coverage, Consistency, Coherence, Fluency, etc.	ChatGPT
[80]	2023	Coherence, Consistency, Fluency, Relevance	GPT-3
[128]	2023	Coverage, Factuality, Coherence, Informativeness, etc.	GPT-4
[220]	2023	General	GPT-3
[221]	2023	Consistency, Relevance, Faithfulness	GPT-4
[187]	2023	Coherence, Consistency, Fluency, Relevance	GPT-4, ChatGPT
[135]	2023	Consistency	ChatGPT
[82]	2023	Consistency	LLaMa
[26]	2023	Consistency	Multiple
[232]	2024	Consistency	Llama-2, Vicuna, ChatGPT
[30]	2024	Meta-evaluation	GPT-3.5, Claude, Gemini, etc.

the balance between performance and resource requirements. Additionally, Fu et al. [58] investigate the viability of smaller, more compact LLMs as alternatives to their larger counterparts. They find that most smaller LLMs, even after fine-tuning, did not surpass larger zero-shot LLMs in summarization tasks.

In addition, the evaluation of LLMs in the multilingual setting [213] reveals that ChatGPT and GPT-4 tend to produce detailed, lengthy summaries, and can achieve a better balance between informativeness and conciseness through interactive prompting. Huang et al. [78] explore applying

LLMs in MDS and point out that the task remains a complex challenge for LLMs mainly due to their limited coverage.

4.1.2 Controllable/Query-Focused Summarization. Yang et al. [241] conduct an evaluation of ChatGPT's performance on query-based summaries, revealing that ChatGPT's performance is comparable to traditional fine-tuning methods in terms of ROUGE scores. Liu et al. [127] evaluate LLMs on instruction-controllable text summarization, where the model input consists of both a source article and a natural language requirement for the desired summary characteristics. The study finds that instruction-controllable text summarization remains a challenging task for LLMs due to persistent factual errors in the generated summaries, a lack of strong alignment between LLM-based evaluation methods and human annotators when assessing the quality of candidate summaries, and performance discrepancies among different LLMs.

Additionally, Pu and Demberg [164] conduct a systematic inspection of ChatGPT's performance in adapting its output to different target audiences (experts vs. laymen) and writing styles (formal vs. informal). Their findings indicate that the stylistic variations produced by humans are considerably larger than those demonstrated by ChatGPT. Moreover, ChatGPT sometimes incorporates factual errors or hallucinations when adapting the text to suit a specific style.

4.1.3 LLM Summarization Characteristics. Recent benchmarking studies have also explored the characteristics of LLMs when generating summaries, including generation faithfulness, fairness, and position bias. These studies aim to understand how accurately LLMs produce summaries that reflect the source content and how the position of information within the source text affects the generated summaries.

Faithfulness and Factuality: Tam et al. [197] propose a factual inconsistency benchmark to measure LLM factuality preferences in summarization. The benchmark includes human-verified factual summaries and annotated inconsistent summaries, providing a means to assess a model's factual consistency. The study finds that while existing LLMs generally prefer factually consistent summaries, they often fail to detect inconsistencies when content is copied verbatim.

Tang et al. [198] introduce an evaluation benchmark for topic-focused dialogue summarization, which includes human annotations of factual consistency along with detailed explanations. Their analysis reveals that LLMs often generate significant amounts of factual errors in the dialogue domain and perform poorly as binary factual evaluators. Laban et al. [101] reach similar conclusions, finding that most LLMs struggle to detect factual inconsistency, with performance close to random chance. Additionally, they highlight that some existing competitive results are caused by issues with existing evaluation benchmarks. Overall, there are still gaps in LLMs' ability to reason about facts and detect inconsistencies.

Fairness: Zhang et al. [264] define summary fairness as not underrepresenting perspectives of any groups of people and find that both the model-generated and the human-written reference summaries suffer from low fairness.

Position bias: Lead bias is a common phenomenon in news summarization, where the early parts of an article often contain the most salient information. Ravaut et al. [171] conduct a comprehensive study on context utilization and position bias in summarization with LLMs, revealing that models tend to favor the initial and final segments of the input but largely ignore the middle, resulting in a U-shaped performance pattern. Similarly, Chhabra et al. [31] examine zero-shot abstractive summarization by measuring position bias, which refers to the tendency of a model to unfairly prioritize information from certain parts of the input text over others. Their study shows that LLMs exhibit significant lead bias in extreme summarization tasks.

4.1.4 Applications. Recent research also explores applying LLMs in interdisciplinary domains like medical document processing and source code.

Medical Document Summarization: Van Veen et al. [204] use LLMs to summarize clinical documents, including radiology reports, patient questions, progress notes, and doctor-patient dialogue. Their clinical reader study finds that summaries from the best-adapted LLMs are deemed either equivalent or superior compared to those from medical experts. However, these LLM-generated summaries face safety challenges such as hallucinations and factual errors, which could potentially lead to medical harm.

In contrast, Tang et al. [199] apply LLMs in zero-shot medical evidence summarization across different clinical domains. They demonstrate that LLMs are susceptible to generating factually inconsistent summaries and making overly convincing or uncertain statements, which could lead to misinformation. Additionally, the models struggle to identify salient information and are more error-prone when summarizing longer texts.

Shaib et al. [183] enlist domain experts to evaluate summaries of biomedical articles generated by GPT-3 in both single- and multi-document settings. The study finds that while GPT-3 could faithfully summarize and simplify single biomedical articles, it struggled to provide accurate aggregations of findings across multiple documents.

Code Summarization: Ahmed and Devanbu [3] first examine the application of few-shot training with GPT Codex and demonstrate its ability to outperform state-of-the-art models for code summarization through the utilization of project-specific training. Subsequently, Sun et al. [194] assess ChatGPT's performance on the extensively employed Python dataset, CSN-Python, revealing that its code summarization performance falls short in comparison to contemporary models in terms of BLEU and ROUGE-L metrics.

Moreover, Luo et al. [135] delve into the potential of LLMs in comprehending binary code. They discover that LLMs exhibit superior performance when applied to decompiled code in comparison to IR code, assembly code, and raw bytes. Notably, they identify that function names play a significant role in shaping the semantics of decompiled code. Furthermore, their research reveals that zero-shot prompts surpass few-shot and **chain-of-thought (CoT)** prompts both in terms of performance and cost efficiency for large-scale tasks. Similarly, Halder and Hockenmaier [71] observe that the performance of models on individual instances often hinges on the subword token overlap between the code and corresponding natural language descriptions, particularly prevalent in function names.

4.2 Modeling Studies

LLM-based summarization differs from prior methods by shifting focus from task-specific model development to leveraging large-scale PLMs in zero-shot and few-shot settings. While earlier approaches primarily involved designing novel architectures and optimizing models on standard datasets, LLM-based techniques emphasize prompting, multi-agent collaboration, and alignment strategies to improve summary quality. Additionally, distillation and specialized training methods enhance the efficiency of LLMs for summarization. Beyond serving as standalone summarization models, LLMs are also increasingly employed as auxiliary tools to refine and enhance summarization systems, marking a fundamental shift in research priorities and methodologies.

As reviewed in Section 3, significant progress in text summarization has been achieved over the past decade, largely due to the paradigm shift toward deep learning methods and fine-tuning PLMs. Most previous summarization research has focused on developing new models, algorithms, systems, or architectures and validating their effectiveness on standard datasets, referred to as modeling studies. Recent advancements in summarization research also introduce new modeling approaches with LLMs, primarily in zero-shot and few-shot settings. These studies encompass a variety of methods, including prompting-based, multi-agent-based, and alignment-based techniques to enhance the quality of generated summaries. Additionally, distillation-based methods and

innovative training strategies have been proposed to improve model efficiency. Furthermore, LLMs are increasingly utilized as tools to enhance system performance in text summarization tasks.

4.2.1 Prompting-Based Method. PromptSum [172] proposes a method that combines prompt tuning with a multi-task objective and discrete entity prompts for abstractive summarization. This approach employs a soft prompt to generate an entity chain, which is then used alongside another soft prompt for summary generation, thereby enhancing model efficiency and controllability. Subsequently, Bhaskar et al. [10] investigate prompting pipeline methods for long-form opinion summarization, including hierarchical summarization with chunking and pre-extraction using an extractive summarization model through a topic clustering-chunking-generation pipeline. Additionally, Chang et al. [22] introduce a prompting method to hierarchically merge chunk-level summaries and incrementally update a running summary, focusing on the application of LLMs for book-length summaries. Sun et al. [192] study prompt chaining that performs drafting, critiquing, and refining phases through a series of three discrete prompts versus a stepwise prompt that integrates these phases within a single prompt.

Chain of Thought: CoT methods have demonstrated significant potential in eliciting the reasoning capabilities of LLMs, thereby enhancing the quality of generated summaries [216]. Building on this, SumCoT [215] expands CoT into an element-aware summarization approach that prompts the LLM to first enumerate important facts and subsequently integrate these facts into a coherent summary, ensuring that the final output is both comprehensive and well-organized. Furthermore, **chain-of-density (CoD)** [1] iteratively incorporates missing salient entities into initially entity-sparse summaries while maintaining a fixed length, resulting in more informative summaries with reduced lead bias. Additionally, ChartThinker [120] proposes to synthesize deep analysis based on CoT and strategies of context retrieval, aiming to improve the logical coherence and accuracy of the generated summaries.

4.2.2 Multi-Agent Based Method. Recent research has explored the construction of multi-agent systems utilizing LLMs to enhance summarization quality. SummIt [255] introduces an iterative text summarization framework based on ChatGPT, incorporating multiple LLM agents such as summarizers and evaluators. This system refines the generated summaries through a process of self-evaluation and feedback, and also investigates the benefits of integrating knowledge and topic extractors to improve the faithfulness and controllability of the summaries. Similarly, ImpressionGPT [137] designs an iterative optimization algorithm to perform automatic evaluations of the generated impressions and formulates corresponding instruction prompts to further refine the model's output. This approach leverages the ICL capabilities of LLMs by creating dynamic contexts with domain-specific, individualized data.

Moreover, ISQA [112] proposes iterative factuality refining for **informative scientific question-answering (ISQA)** feedback. It operates in a detailed manner, instructing the summarization agent to reinforce validated statements from positive feedback and correct inaccuracies highlighted in negative feedback. SliSum [110] employs sliding windows and self-consistency with LLMs. It divides the source article into overlapping windows, allowing LLMs to generate local summaries for each window's content and then aggregates them using clustering and a majority voting algorithm.

4.2.3 Alignment. Recent advancements in summarization techniques have focused on leveraging feedback mechanisms and instruction-based fine-tuning to improve the quality and user alignment of generated summaries [158]. Liu et al. [126] first investigate the use of informational feedback from human annotations in natural language to enhance generation quality and factual consistency. Subsequently, SALT [243] introduces a method that integrates both human-edited and model-generated data within the training loop by simulating human edits. InstructPTS [56] proposes

using instruction fine-tuning of LLMs to achieve controllable product title summarization across different dimensions, utilizing automatically generated instructions derived from a parallel dataset.

4.2.4 Distillation-Based Method. Despite the strong summarization performance of LLMs, their resource demands have limited their widespread use. Additionally, privacy concerns have arisen when using LLM-as-a-service APIs, particularly for sensitive data. This underscores the need for more compact local models that can still effectively capture summarization abilities.

Liu et al. [129] proposed training smaller models like BART or BRIO using contrastive learning, where LLMs such as ChatGPT serve as evaluators to signal which generated summary candidate is superior. TriSum [84] first extracts a set of aspect-triple rationales and summaries with LLMs, and then trains a smaller local model with a curriculum learning strategy on these tasks. Triple here is a structure formatting a piece of free text into a subject, a relation, and an object.

Additionally, InheritSumm [235] achieves a compact summarization model with versatile capabilities across different settings by training using knowledge distillation from the GPT-3 model, mimicking the GPT-generated summaries on general documents. Similarly, Jung et al. [87] present Impossible Distillation, leveraging the paraphrastic proximity inherent in PLMs like GPT-2, identifying and distilling high-quality paraphrases from these subspaces.

4.2.5 Training Strategy. Recent research also explores training strategies to fine-tune LLMs. For example, Xia et al. [222] propose a method to alleviate hallucinations in LLM outputs using hallucination diversity-aware active learning. This approach selects diverse hallucinations for annotations in the active learning process for LLM fine-tuning.

Whitehouse et al. [217] investigate the potential of **parameter-efficient fine-tuning (PEFT)** with low-rank adaptation (LoRA) [77] in the domain of multilingual summarization. Furthermore, Pu and Demberg [165] propose to integrate RST with LoRA to enhance long document summarization performance.

4.2.6 Others. Recent research efforts have also explored innovative approaches to leverage LLMs as a tool for summarization tasks, such as data cleaning for QFS summarization in [104], data annotation for pre-training for dialogue summarization in [276], and pseudo-label generation for dialogs to fine-tune smaller model in [145].

4.3 Evaluation Studies

Evaluating the quality of generated summaries accurately has indeed been a major challenge in summarization research. As discussed in Section 2.2, researchers have explored reference-based and reference-free automatic metrics, but the correlation with human judgment is still limited and cannot reliably evaluate summary quality [55]. With the recent emergence of LLMs, researchers are also investigating their use in summary quality evaluation.

4.3.1 General Summary Evaluation. GPTScore [57] represents a pioneering effort in utilizing LLMs for summary evaluation metrics. It leverages the emergent abilities of LLMs by incorporating task specifications and detailed aspect definitions, which serve as instructions for scoring generated texts. Building on this, Gao et al. [61] investigate ChatGPT's capacity for human-like summarization evaluation. Their findings indicate that ChatGPT can perform annotations relatively smoothly using various methods, including Likert scale scoring, pairwise comparison, the Pyramid method, and binary factuality evaluation.

Similarly, Jain et al. [80] propose using LLMs as multi-dimensional evaluators for text summarization through ICL, assessing aspects such as fluency, coherence, and factuality. G-Eval [128], another notable LLM-based evaluation metric for summaries, uses a CoT and form-filling paradigm. G-Eval demonstrates a high Spearman correlation with human evaluations on summarization tasks,

significantly outperforming previous methods. Wu et al. [220] introduce role-players prompting mechanism to evaluate summaries from both objective and subjective perspectives. Another approach proposed by Wu et al. [221] aims to reduce evaluation costs with an Extract-then-Evaluate method.

Despite these advancements, Shen et al. [187] investigate the stability and reliability of LLMs as automatic evaluators for abstractive summarization. They find that although models like ChatGPT and GPT-4 outperform traditional automatic metrics, they are not yet reliable substitutes for human evaluators. LLM evaluators exhibit inconsistent ratings across different candidate systems and dimensions, struggle with close performance comparisons, and show a lower correlation with human judgments as the quality of summaries improves.

4.3.2 Factual Consistency Evaluation. Luo et al. [135] first explore ChatGPT's ability to evaluate factual inconsistency under a zero-shot setting. They find that ChatGPT achieves satisfying performance but prefers more lexically similar candidates and instances of false reasoning. Chen et al. [26] further extend the setting with different LLMs and different prompting techniques.

Subsequently, Jia et al. [82] propose a zero-shot faithfulness evaluation with a foundation language model without prompting or instruction tuning. Xu et al. [232] propose three zero-shot paradigms for factual consistency evaluation, including direct inference on the entire summary or each summary window, and entity verification through question generation and answering.

Compared to traditional entailment-based or QA-based factual consistency evaluation, LLM-based methods benefit from extensive pre-training on large corpora and enhanced verification through information retrieval. As LLMs become central to summarization, factual consistency evaluation must be tailored to different domains. In news summarization, evaluation metrics should prioritize detecting misinformation, bias, and source credibility. For scientific and academic summarization, the focus should be on verifying citation accuracy and ensuring alignment with original claims. Domain-specific adaptations are essential to maintaining reliability.

4.3.3 Others. Chang et al. [22] propose an automatic metric BOOOOKSCORE to assess summary coherence in book-length summary generation. Huang et al. [78] explore the usage of automatic LLM-based metrics under the MDS setting. In addition, Chern et al. [30] conduct a meta-evaluation that can effectively, reliably, and efficiently evaluate the performance of LLMs as evaluators across diverse tasks and scenarios via agent debate.

5 Open Problems and Future Directions

This section discusses the trends in recent LLM-based summarization, the general challenges of text summarization that have yet to be solved and pinpoints potential future research directions to attract practitioners' attention and improve our understanding and techniques in text summarization in the new era of LLMs.

5.1 Quantitative Results

The CNN/DM dataset [73] has been used as the predominant benchmark for assessing summary quality and comparing the efficacy of summarization methodologies in recent years. Here, we also review the quantitative results of representative summarization approaches on it, as delineated in Table 6, to glean insights into the trajectory of research advancements. Despite the limitations discussed in Section 2.2, ROUGE F-scores remain the standard and most widely adopted way for the assessment of summary quality, so we report ROUGE scores in the table for comprehensive evaluation.

The quantitative results in Table 6 reveal significant advancements in summarization systems over recent years, driven by paradigm shifts toward deep learning and PLM-based approaches.

Table 6. Quantitative Results on the CNN/DM Dataset

Method	Model	Year	ROUGE-1	ROUGE-2	ROUGE-L
Extractive(unsup.)	TextRank[143]	2004	33.2	11.8	29.6
Extractive	SUMMARUNNER [150]	2017	39.60	16.20	35.30
Extractive	REFRESH [153]	2018	40.00	18.20	36.60
Extractive	BANDITSUM [48]	2018	41.50	18.70	37.60
Extractive	NEUSUM [275]	2018	41.59	19.01	37.98
Extractive (unsup.)	PACSUM [269]	2019	40.7	17.8	36.9
Extractive	HIBERT [263]	2019	42.37	19.95	38.83
Extractive	BERTSUMEXT [131]	2019	43.85	20.34	39.90
Extractive	DISCOBERT [231]	2019	43.77	20.85	40.67
Extractive	HSG [211]	2020	42.95	19.76	39.23
Extractive	MATCHSUM [271]	2020	44.41	20.86	40.55
Extractive	DiffuSum [253]	2023	44.83	22.56	40.56
Abstractive	PTGEN [181]	2017	36.44	15.66	33.42
Hybrid	BOTTOMUP [63]	2018	41.22	18.68	38.34
Abstractive	DCA [21]	2018	41.69	19.47	37.92
Abstractive	BERTSUMABS [131]	2019	41.72	19.39	38.76
Abstractive	BART [107]	2019	44.16	21.28	40.90
Abstractive	PEGASUS [258]	2020	44.17	21.47	41.11
Hybrid	GSum [49]	2020	45.94	22.32	42.48
Abstractive	SimCLS [132]	2021	46.67	22.15	43.54
Abstractive	SummIt (ChatGPT) [255]	2023	37.29	13.60	26.87
Abstractive	Element (GPT-3) [215]	2023	37.75	15.20	34.25
Abstractive	SliSum (Claude2) [110]	2024	47.75	23.16	44.26
Abstractive	HADAS [222]	2024	-	-	20.12

Results are taken from the original paper if available.

LLM-based methods achieve competitive ROUGE scores, while offering benefits such as generalization, instruction-following, and zero-shot capability. Despite CNN/DM's establishment as a predominantly extractive dataset [270], abstractive summarization systems have demonstrated superior flexibility and performance compared to extractive methods. However, it's worth noting that ROUGE scores may inadequately capture summary quality, particularly in the new era of LLMs and when reference summaries are of low quality, as discussed in [66, 261]. Consequently, recent research has shifted focus toward human evaluations to more accurately assess summarization system performance in the context of evolving language models.

Selecting the appropriate summarization approach depends on the specific task requirements, domain constraints, and computational resources available. For extractive summarization, traditional statistical methods like TextRank [143] remain viable for applications where efficiency is prioritized. Deep learning-based extractive models, such as BERTSUM [131], are more suitable when sentence-level coherence and contextual relevance are critical, particularly in news summarization and legal document processing. Abstractive summarization, powered by PLMs is ideal for generating more natural and human-like summaries, making it well-suited for conversational AI, customer support, and content creation. However, when dealing with domain-specific content—such as medical or scientific summarization, LLMs can leverage their broad knowledge and few-shot capabilities to generate high-quality summaries without extensive fine-tuning. Yet, LLM-based approaches may struggle with factual consistency and hallucination, necessitating additional verification techniques

such as QA-based evaluation [50] and **retrieval-augmented generation (RAG)**. For resource-constrained settings, parameter-efficient tuning methods like LoRA [77] enable adaptation of LLMs without incurring the full cost of fine-tuning. Ultimately, the choice of summarization method should be guided by the balance between quality, efficiency, and domain specificity, with hybrid approaches increasingly emerging as a promising solution to maximize performance across different tasks.

5.2 Computational Efficiency

Computational efficiency has been a critical factor in the evolution of summarization methods, with each paradigm shift bringing changes in resource demands and scalability. Early statistical approaches, such as extractive methods based on TF-IDF or graph-based algorithms like TextRank [143], were computationally lightweight and required minimal training, making them suitable for low-resource settings. However, these methods lacked contextual understanding and struggled with coherence. The introduction of deep learning and sequence-to-sequence models, such as RNNs and early transformer-based models like BERTSUM [131], significantly improved summarization quality but required extensive supervised training on large datasets, increasing computational costs. The rise of PLMs, including BART and T5 [107], further advanced summarization by leveraging transfer learning, reducing the need for task-specific training while still requiring fine-tuning, which demands substantial GPU resources. With the advent of LLMs such as GPT-4, summarization has become more flexible, supporting zero-shot and few-shot inference without explicit fine-tuning. However, LLMs come with immense computational costs due to their billions of parameters, requiring powerful hardware for both training and inference. Techniques such as parameter-efficient tuning [77] (e.g., LoRA and adapter layers) and distillation have been explored to mitigate these costs, yet efficiency remains a key challenge, particularly for real-time applications and domain-specific summarization. Future research must balance performance and efficiency, exploring optimized inference strategies, model compression, and hardware-aware deployment to make LLM-based summarization more accessible and scalable [42, 70, 234].

5.3 Open Challenges

Despite the significant progress made in recent years, there are still open challenges in text summarization research in the new era of LLMs.

Hallucination: Hallucination is one of the most significant issues hindering LLMs from real-world deployments. It refers to instances where the model generates information that is factually incorrect or not present in the source text, including fabricating facts, events, or details that were never mentioned in the original document. This issue is particularly problematic in critical fields such as medicine and law, where accuracy is paramount. Researchers have explored various approaches to measure and quantify the level of hallucination through benchmarks [118], and to mitigate hallucination in LLM generation with enhanced training protocols [81, 265], RAG [24, 108, 113, 246], self-reflection [188, 255], and post-processing techniques [23, 25]. However, more efforts are still needed to achieve faithful and factual summary generation.

Bias: LLMs can reflect and amplify biases present in their training data, leading to the generation of biased summaries [60, 154, 264]. Addressing these ethical issues involves developing methods to detect bias in generation and mitigate biased content in the summary generation process. There is a growing need for models that produce fair and unbiased summaries, particularly in socially sensitive domains.

Computational Efficiency: The computational efficiency of LLMs is a critical factor influencing their performance, scalability, and practical application. These models require significant computational resources due to their extensive parameter sizes and the complexity of training and

inference. Efficiency improvements can be achieved through hardware acceleration [239, 240] and techniques like model pruning [138, 191], quantization [111, 223], PEFT [42, 77], and knowledge distillation [70, 87, 234]. Despite these advancements, balancing computational efficiency with model accuracy and generalization remains a challenging and ongoing area of research.

Personalization: Creating summaries personalized to individual user preferences poses another significant challenge [93, 219]. This task involves understanding user interests, reading habits, and prior knowledge to generate summaries that are relevant and useful to the specific user. It requires NLP techniques that not only condense information but also align it with the user's unique profile. Personalized summarization can be particularly useful in fields such as education, news, and media, where it can adapt content to different learning styles and present information that aligns with the reader's interests [14, 29, 163]. However, this personalization must be balanced with privacy considerations, ensuring that user data is handled securely and ethically.

Interpretability and Explainability: Interpretability and explainability of LLM-based summarization are crucial for building trust and ensuring transparency in AI-generated content [136]. Interpretability refers to the ability to understand how a model makes its decisions, while explainability involves elucidating the reasons behind specific outputs. In the context of summarization, these aspects help users comprehend why certain information was highlighted or omitted [109, 178]. Providing clear explanations for summarization outputs is particularly important in sensitive domains like healthcare and legal fields, where understanding the rationale behind the summary can impact decision-making processes [27]. Despite advancements, achieving high levels of interpretability and explainability remains challenging due to the complexity and black-box nature of LLMs.

5.4 Future Directions

Here, we also highlight some future directions for summarization in the era of LLMs to broaden the scope of existing summarization systems, thereby expanding their application scenarios and impacts.

Summarization beyond Text: Over the past few decades, significant progress has been made in the field of text summarization research. The advent of powerful LLMs such as GPT-4 and GPT-4o [157], capable of processing various data formats, has catalyzed advancements in multimodal summarization. This emerging area extends traditional text summarization to include other forms of data such as images [83, 147, 168, 224, 279, 280], tables [122], source code [209, 214], audio [88], and video [97, 117]. Multimodal summarization systems aim to provide more comprehensive and contextually rich summaries by integrating diverse data sources. These systems enable the generation of summaries that incorporate key visual elements from videos or important audio clips from podcasts, thereby enhancing user understanding and expanding real-world applications. This holistic approach has the potential to revolutionize information consumption, making it more accessible and engaging across various formats, and represents a significant step toward the development of AI assistants.

New Tasks in Summarization: Traditional summarization research has heavily focused on generic and extractive approaches. However, with the advent of powerful LLMs, the horizon of text summarization is brimming with the promise of tackling novel tasks that extend beyond conventional settings. Emerging research is steering toward more nuanced aspects tailored to specific user scenarios and real-world applications.

These include personalized and customized summarization that aligns with individual preferences, interests, and reading history [29, 274], human-in-the-loop summarization enabling user engagement and verification for high-stakes domains like legal and medical fields [4, 184, 185], real-time summarization that enhances the speed and efficiency of models for live events or breaking

news [33, 182, 238, 266], sentiment-aware summarization reflecting the underlying sentiment of source material [230, 237], and diversified summarization consolidating information across multiple articles [78], and beyond. These new frontiers not only broaden the scope of summarization tasks but also hold immense potential in advancing the sophistication and utility of LLM-generated summaries across various domains and applications.

Ethical and Responsible Studies: Developing ethical AI practices is crucial for the future of summarization [60, 114]. This includes creating transparent models that can explain their summarization processes, ensuring fairness in summaries that do not underrepresent perspectives of any groups of people, and actively working to mitigate biases [76, 264]. There is also an urgent need for more theoretical groundings and ethical considerations for model generation from cognitive science, social science, psychology, and linguistics. Effectively and safely incorporating updated world knowledge and common sense into the summary generation cycle with control is also important to ensure the summaries are accurate, up-to-date, and factually correct [25]. Ethical considerations will be paramount in building trust and ensuring the responsible use of summarization technologies.

Domain Specific Summarization LLM: Current LLMs demonstrate extraordinary generalization capability across different domains and subjects, thanks to large-scale pre-training with web data. However, this generalization capability may limit their practical use in specific domains when operating in the zero-shot setting due to the lack of domain-specific knowledge, rules, patterns, and terminologies. Consequently, the need for domain-specific LLM-based summarization systems is growing and necessary for practical deployment [40, 89, 202, 229]. Effectively and efficiently adapting open-sourced LLMs to domain-specific applications is an important direction in building trustworthy expert summarization systems.

6 Conclusion

In conclusion, the landscape of text summarization research has witnessed profound advancements driven by deep neural networks, PLMs, and the recent emergence of LLMs. This survey captures the evolution of text summarization paradigms in detail, offering a comprehensive examination of both pre-LLM and LLM eras. Furthermore, the survey highlights research trends and delineates open challenges, proposing future research directions crucial for advancing summarization techniques. Through this comprehensive synthesis, the survey aims to serve as a valuable resource for researchers and enthusiasts, facilitating informed exploration and innovation in the ever-evolving domain of text summarization.

References

- [1] Griffin Adams, Alexander Fabbri, Faisal Ladhak, Eric Lehman, and Noémie Elhadad. 2023. From sparse to dense: GPT-4 summarization with chain of density prompting. In *Proceedings of the 4th New Frontiers in Summarization Workshop*. Singapore. Association for Computational Linguistics, 6874.
- [2] Armen Aghajanyan, Akshat Shrivastava, Anchit Gupta, Naman Goyal, Luke Zettlemoyer, and Sonal Gupta. 2020. Better fine-tuning by reducing representational collapse. *arXiv preprint arXiv:2008.03156*.
- [3] Toufique Ahmed and Premkumar Devanbu. 2022. Few-shot training LLMs for project-specific code-summarization. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*. 1–5.
- [4] P. V. S. Avinesh, Benjamin Hättasch, Orkan Özyurt, Carsten Binnig, and Christian M. Meyer. 2018. Sherlock: A system for interactive summarization of large text collections. *Proceedings of the VLDB Endowment* 11, 12 (August 2018), 1902–1905.
- [5] Sanghwan Bae, Taeuk Kim, Jihoon Kim, and Sang-goo Lee. 2019. Summary level training of sentence rewriting for abstractive summarization. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*. Hong Kong, China. Association for Computational Linguistics, 10–20.
- [6] Seyed Ali Bahrainian, Sheridan Feucht, and Carsten Eickhoff. 2022. NEWTS: A corpus for news topic-focused summarization. In *Findings of the Association for Computational Linguistics: ACL 2022*. Dublin, Ireland. Association for Computational Linguistics, 493–503.

- [7] Tal Baumel, Raphael Cohen, and Michael Elhadad. 2016. Topic concentration in query focused summarization datasets. In *Proceedings of the AAAI Conference on Artificial Intelligence* 30, 1 (2016), 2573–2579.
- [8] Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. arXiv:2004.05150. Retrieved from <https://arxiv.org/abs/2004.05150>
- [9] Anja Belz and Albert Gatt. 2008. Intrinsic vs. extrinsic evaluation measures for referring expression generation. In *Proceedings of ACL-08: HLT*. 197–200.
- [10] Adithya Bhaskar, Alex Fabbri, and Greg Durrett. 2023. Prompted opinion summarization with GPT-3.5. In *Findings of the Association for Computational Linguistics: ACL 2023*. Toronto, Canada. Association for Computational Linguistics, 9282–9300.
- [11] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research* 3 (January 2003), 993–1022.
- [12] Sergey Brin and Lawrence Page. 1998. The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems* 30, 1–7 (1998), 107–117.
- [13] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems* 33 (2020), 1877–1901.
- [14] Pengshan Cai, Kaiqiang Song, Sangwoo Cho, Hongwei Wang, Xiaoyang Wang, Hong Yu, Fei Liu, and Dong Yu. 2023. Generating user-engaging news headlines. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Toronto, Canada. Association for Computational Linguistics, 3265–3280.
- [15] Dimas O. Cajueiro, Ana G. Nery, Iuri Tavares, M. K. de Melo, S. A. dos Reis, L. Weigang, and V. R. Celestino. 2024. Comprehensive review of automatic text summarization techniques. *Computing & Informatics* 43, 5 (2024).
- [16] Meng Cao. 2022. A survey on neural abstractive summarization methods and factual consistency of summarization. arXiv:2204.09519. Retrieved from <https://arxiv.org/abs/2204.09519>
- [17] Meng Cao, Yue Dong, Jiapeng Wu, and Jackie Chi Kit Cheung. 2020. Factual error correction for abstractive summarization models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online. Association for Computational Linguistics, 6251–6258.
- [18] Shuyang Cao and Lu Wang. 2021. CLIFF: Contrastive learning for improving faithfulness and factuality in abstractive summarization. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 6633–6649.
- [19] Ziqiang Cao, Wenjie Li, Sujian Li, Furu Wei, and Yangqiu Li. 2016. AttSum: Joint learning of focusing and summarization with neural attention. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, Osaka, Japan. 547–556.
- [20] Jaime Carbonell and Jade Goldstein. 1998. The use of MMR, diversity-based reranking for reordering documents and producing summaries. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. 335–336.
- [21] Asli Celikyilmaz, Antoine Bosselut, Xiaodong He, and Yejin Choi. 2018. Deep communicating agents for abstractive summarization. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. New Orleans, Louisiana. Association for Computational Linguistics. 1662–1675.
- [22] Yapei Chang, Kyle Lo, Tanya Goyal, and Mohit Iyyer. 2023. BoookScore: A systematic exploration of book-length summarization in the era of LLMs. In *International Conference on Learning Representations (ICLR 2024)*.
- [23] Anthony Chen, Panupong Pasupat, Sameer Singh, Hyung Won Lee, and Kelvin Guu. 2023. PURR: Efficiently editing language model hallucinations by denoising language Model Corruptions. *arXiv preprint arXiv:2305.14908*.
- [24] Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. 2024. Benchmarking large language models in retrieval-augmented generation. In *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 16 (2024), 17754–17762.
- [25] Sihao Chen, Fan Zhang, Kazuo Sone, and Dan Roth. 2021. Improving faithfulness in abstractive summarization with contrast candidate generation and selection. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Online. Association for Computational Linguistics, 5935–5941.
- [26] Shiqi Chen, Siyang Gao, and Junxian He. 2023. Evaluating factual consistency of summaries with large language models. *arXiv preprint arXiv:2305.14069*.
- [27] Zhiyu Zoey Chen, Jing Ma, Xiaoyu Zhang, Ning Hao, An Yan, Armineh Nourbakhsh, Xiaoyang Yang, Julian McAuley, Linda Petzold, and William Yang Wang. 2024. A survey on large language models for critical societal domains: Finance, healthcare, and law. *arXiv preprint arXiv:2405.01769*.
- [28] Jianpeng Cheng and Mirella Lapata. 2016. Neural summarization by extracting sentences and words. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*. 484–494.

- [29] Xiaocheng Cheng, Shen Gao, Yixin Zhang, Yibo Wang, Xiaojie Chen, Minlie Li, Dongyan Zhao, and Rui Yan. 2023. Towards personalized review summarization by modeling historical reviews from customer and product separately. *arXiv preprint arXiv:2301.11682*.
- [30] Shengyu Chern, Ethan Chern, Graham Neubig, and Pengfei Liu. 2024. Can large language models be trusted for evaluation? Scalable meta-evaluation of LLMs as evaluators via agent debate. *arXiv preprint arXiv:2401.16788*.
- [31] Anshuman Chhabra, Hadi Askari, and Prasant Mohapatra. 2024. Revisiting zero-shot abstractive summarization in the era of large language models from the perspective of position bias. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*. 1–11.
- [32] Sangwoo Cho, Logan Lebanoff, Hassan Foroosh, and Fei Liu. 2019. Improving the similarity measure of determinantal point processes for extractive multi-document summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy. Association for Computational Linguistics, 1027–1038.
- [33] S. Cho, F. Dernoncourt, T. Ganter, T. Bui, N. Lipka, W. Chang, H. Jin, J. Brandt, H. Foroosh, and F. Liu. 2021. StreamHover: Livestream transcript summarization and annotation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 6457–6474.
- [34] Sangwoo Cho, Franck Dernoncourt, Tim Ganter, Trung Bui, Nedim Lipka, Walter Chang, Hailin Jin, Jonathan Brandt, Hassan Foroosh, and Fei Liu. 2021. StreamHover: Livestream transcript summarization and annotation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Online and Punta Cana, Dominican Republic. 6457–6474.
- [35] Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems* 30 (2017).
- [36] Eric Chu and Peter J. Liu. 2018. Unsupervised neural multi-document abstractive summarization. arXiv:1810.05739. Retrieved from <https://arxiv.org/abs/1810.05739>
- [37] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, et al. 2024. Scaling Instruction-Finetuned Language Models. *Journal of Machine Learning Research* 25, 70 (2024), 1–53.
- [38] Elizabeth Clark, Asli Celikyilmaz, and Noah A. Smith. 2019. Sentence Mover’s similarity: Automatic evaluation for multi-sentence texts. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy, 2748–2760.
- [39] Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. A discourse-aware attention model for abstractive summarization of long documents. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, New Orleans, Louisiana, 615–621.
- [40] Jiayi Cui, Zongjian Li, Yang Yan, Bohua Chen, and Li Yuan. 2023. ChatLaw: Open-source legal large language model with integrated external knowledge bases. *CoRR*, abs/2306.16092.
- [41] Hoa Trang Dang and Karolina Owczarzak. 2008. Overview of the TAC 2008 update summarization task. In *Text Analysis Conference (TAC)*.
- [42] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. QLoRA: Efficient finetuning of quantized LLMs. In *Advances in Neural Information Processing Systems* 36 (2023), 10088–10115.
- [43] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Minneapolis, Minnesota. Association for Computational Linguistics, 4171–4186.
- [44] Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. In *Advances in Neural Information Processing Systems* 32 (2019).
- [45] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, Xu Sun, Lei Li, and Zhifang Sui. 2024. A survey on in-context learning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, Miami, Florida, USA. Association for Computational Linguistics, 1107–1128.
- [46] Yue Dong. 2018. A survey on neural network-based summarization methods. arXiv:1804.04589. Retrieved from <https://arxiv.org/abs/1804.04589>
- [47] Yue Dong, Andrei Mircea, and Jackie C. K. Cheung. 2021. Discourse-aware unsupervised summarization for long scientific documents. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*. Online. Association for Computational Linguistics, 1089–1102.
- [48] Yue Dong, Yikang Shen, Eric Crawford, Herke van Hoof, and Jackie Chi Kit Cheung. 2018. BanditSum: Extractive summarization as a contextual bandit. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun’ichi Tsujii (Eds.), Association for Computational Linguistics, 3739–3748. DOI: <https://doi.org/10.18653/V1/D18-1409>

- [49] Zi-Yi Dou, Pengfei Liu, Hiroaki Hayashi, Zhengbao Jiang, and Graham Neubig. 2021. GSum: A general framework for guided neural abstractive summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Online, 4830–4842.
- [50] Esin Durmus, He He, and Mona Diab. 2020. FEQA: A question answering evaluation framework for faithfulness assessment in abstractive summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online. Association for Computational Linguistics, 5055–5070.
- [51] Harold P. Edmundson. 1969. New methods in automatic extracting. *Journal of the ACM (JACM)* 16, 2 (1969), 264–285.
- [52] Wael Salah El-Kassas, Cherif Rahhal Salama, Ahmed Abdou Rafea, and Hoda Kamel Mohamed. 2021. Automatic text summarization: A comprehensive survey. *Expert Systems with Applications* 165 (2021), 113679.
- [53] Günes Erkan and Dragomir R. Radev. 2004. Lexrank: Graph-based lexical centrality as salience in text summarization. *Journal of Artificial Intelligence Research* 22 (2004), 457–479.
- [54] Alexander Richard Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. Multi-News: A large-scale multi-document summarization dataset and abstractive hierarchical model. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 1074–1084.
- [55] Alexander R. Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. Summeval: Re-evaluating Summarization Evaluation. *Transactions of the Association for Computational Linguistics* 9 (2021), 391–409.
- [56] Besnik Fetahu, Zhiyu Chen, Oleg Rokhlenko, and Shervin Malmasi. 2023. InstructPTS: Instruction-Tuning LLMs for product title Summarization. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track*. Singapore. Association for Computational Linguistics, 663–674.
- [57] Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2024. GPTScore: Evaluate as you desire. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*. Mexico City, Mexico. Association for Computational Linguistics, 6556–6576.
- [58] Xue-Yong Fu, Md Tahmid Rahman Laskar, Elena Khasanova, Cheng Chen, and Shashi Bhushan TN. 2024. Tiny titans: Can smaller large language models punch above their weight in the real world for meeting summarization? In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 6: Industry Track)*. Mexico City, Mexico, 387–394.
- [59] Yue Fu, Deyi Xiong, and Yue Dong. 2023. Inverse reinforcement learning for text summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore*. Association for Computational Linguistics, 6559–6570.
- [60] Isabel O. Gallegos, Ryan A. Rossi, Jason Barrow, Md Mahfuzur Rahman Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruijia Zhang, and Nesreen K. Ahmed. 2024. Bias and fairness in large language models: A survey. *Computational Linguistics* 50, 3 (2024).
- [61] Mingqi Gao, Jie Ruan, Renliang Sun, Xunjian Yin, Shiping Yang, and Xiaojun Wan. 2023. Human-like Summarization Evaluation with ChatGPT. *arXiv preprint arXiv:2304.02554*.
- [62] Yang Gao, Wei Zhao, and Steffen Eger. 2020. SUPERT: Towards new Frontiers in unsupervised evaluation metrics for multi-document summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 1347–1354.
- [63] Sebastian Gehrmann, Yuntian Deng, and Alexander M. Rush. 2018. Bottom-up abstractive summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels, Belgium. Association for Computational Linguistics, 4098–4109.
- [64] Demian Gholipour Ghalandari, Chris Hokamp, Nghia The Pham, John Glover, and Georgiana Ifrim. 2020. A large-scale multi-document summarization dataset from the wikipedia current events portal. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online. Association for Computational Linguistics, 1302–1308.
- [65] Barbara Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. SAMSum Corpus: A human-annotated dialogue dataset for abstractive summarization. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*. Hong Kong, China. Association for Computational Linguistics, 70–79.
- [66] Tanya Goyal, Junyi Li, and Greg Durrett. 2022. News summarization and evaluation in the era of GPT-3. *arXiv preprint arXiv:2209.12356*.
- [67] Tanya Goyal, Junyi Li, and Greg Durrett. 2022. SNaC: Coherence error detection for narrative summarization. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022)*. Abu Dhabi, United Arab Emirates. Association for Computational Linguistics, 444–463.
- [68] Tanya Goyal and Greg Durrett. 2020. Evaluating factuality in generation with dependency-level entailment. In *Findings of the Association for Computational Linguistics: EMNLP 2020*.
- [69] Matthew Grusky, Matan Naaman, and Yoav Artzi. 2018. Newsroom: A dataset of 1.3 Million summaries with diverse extractive strategies. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. 708–719.

- [70] Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. 2023. MiniLLM: Knowledge distillation of large language models. *arXiv preprint arXiv:2306.08543*.
- [71] Rajarshi Haldar and Julia Hockenmaier. 2024. Analyzing the performance of large language models on code summarization. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING'24)*. 995–1008.
- [72] Tahmid Hasan, Abhik Bhattacharjee, Md Saiful Islam, Kazi Mubasshir, Yi-Fang Li, Young-Bum Kim, M. Sohel Rahman, and Rifat Shahriyar. 2021. XL-Sum: Large-scale multilingual abstractive summarization for 44 Languages. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*. 4693–4703.
- [73] Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems* 28 (2015), 1693–1701.
- [74] Tsutomu Hirao, Yasuhisa Yoshida, Masaaki Nishino, Norihito Yasuda, and Masaaki Nagata. 2013. Single-document summarization as a tree knapsack problem. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP 2013)*. Seattle, Washington, USA. Association for Computational Linguistics, 1515–1520.
- [75] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9, 8 (1997), 1735–1780.
- [76] Chao-Chun Hsu and Chenhao Tan. 2021. Decision-focused summarization. *arXiv:2109.06896*. Retrieved from <https://arxiv.org/abs/2109.06896>
- [77] Edward J. Hu, Yelong Shen, Phil Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *Proceedings of the International Conference on Learning Representations (ICLR)* 1, 2 (2022), 3.
- [78] Kuan-Hao Huang, Peter Laban, Alexander R. Fabbri, Prakhar K. Choubey, Shafiq Joty, Chenyan Xiong, and Caiming Wu. 2024. Embrace Divergence for Richer Insights: A multi-document summarization benchmark and a case study on summarizing diverse information from news articles. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*. 570–593.
- [79] Luyang Huang, Lingfei Wu, and Lu Wang. 2020. Knowledge graph-augmented abstractive summarization with semantic-driven cloze reward. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 5094–5107.
- [80] Shailza Jain, Varun Keshava, Shreyas Muralidhar Sathyendra, Pedro Fernandes, Pengfei Liu, Graham Neubig, and Chunting Zhou. 2023. Multi-dimensional evaluation of text summarization with in-context learning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*. 11327–11347.
- [81] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Comput. Surv.* 55, 12, Article 248 (December 2023), 38 pages. <https://doi.org/10.1145/3571730>
- [82] Qi Jia, Siyu Ren, Yizhu Liu, and Kenny Zhu. 2023. Zero-shot faithfulness evaluation for text summarization with foundation language model. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023)*. Singapore. Association for Computational Linguistics, 11017–11031.
- [83] Chaoya Jiang, Rui Xie, Wei Ye, Jinan Sun, and Shikun Zhang. 2023. Exploiting pseudo image captions for multimodal summarization. In *Findings of the Association for Computational Linguistics: ACL 2023*. Toronto, Canada. Association for Computational Linguistics, 161–175.
- [84] Pengcheng Jiang, Zhihao Zhou, Xinyu Zhang, and Xipeng Qiu. 2024. TriSum: Learning summarization ability from large language models with structured rationale. *arXiv preprint arXiv:2403.10351*.
- [85] Xiang Jiang and Markus Dreyer. 2024. CCSUM: A large-scale and high-quality dataset for abstractive news summarization. (2024).
- [86] Baoyu Jing, Zeyu You, Tao Yang, Wei Fan, and Hanghang Tong. 2021. Multiplex graph neural network for extractive text summarization. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP 2021)*, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics, 133–139.
- [87] Jaehun Jung, Peter West, Liwei Jiang, Faeze Brahma, Ximing Lu, Jillian Fisher, Taylor Sorensen, and Yejin Choi. 2023. Impossible distillation: From low-quality model to high-quality dataset & model for summarization and paraphrasing. *arXiv preprint arXiv:2305.16635*.
- [88] Takatomo Kano, Atsunori Ogawa, Marc Delcroix, Roshan S. Sharma, Kohei Matsuura, and Shinji Watanabe. 2023. Speech summarization of long spoken document: Improving memory efficiency of speech/text encoders. In *Proceedings of the 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2023)*. 1–5.
- [89] Enkelejda Kasneci, Timo K. H. Danks, Julia Koch, Maximilian S. Trautmann, Matthias S. Winkermann, and Wolfgang M. Struckmeier. 2023. ChatGPT for Good? On opportunities and challenges of large language models for education. *Learning and Individual Differences* 103 (2023), 102274. <https://doi.org/10.1016/j.lindif.2023.102274>

- [90] Byeongchang Kim, Hyunwoo Kim, and Gunhee Kim. 2019. Abstractive summarization of reddit posts with multi-level memory networks. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Minneapolis, Minnesota. Association for Computational Linguistics, 2519–2531.
- [91] Yoon Kim. 2014. Convolutional neural networks for sentence classification. arXiv:1408.5882. Retrieved from <https://arxiv.org/abs/1408.5882>
- [92] Thomas N. Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. In *Proceedings of the International Conference on Learning Representations*.
- [93] Hannah Rose Kirk, Scott A. Hale, Bertie Vidgen, and Sandra Wachter. 2024. The benefits, risks and bounds of personalizing the alignment of large language models to individuals. *Nature Machine Intelligence* 6 (2024), 1–10.
- [94] Huan Yee Koh, Jiaxin Ju, Ming Liu, and Shirui Pan. 2022. An empirical survey on long document summarization: Datasets, models, and metrics. *ACM Comput. Surv.* 55, 8, Article 154 (August 2023), 35 pages. <https://doi.org/10.1145/3545176>
- [95] Anastassia Kornilova and Vlad Eidelman. 2019. BillSum: A corpus for automatic summarization of US legislation. *EMNLP-IJCNLP 2019* (2019), 48.
- [96] Mahnaz Koupaei and William Yang Wang. 2018. Wikihow: A large scale text summarization dataset. arXiv:1810.09305. Retrieved from <https://arxiv.org/abs/1810.09305>
- [97] Mateusz Krubiński and Pavel Pecina. 2023. Mlask: Multimodal summarization of video-based news articles. In *Findings of the Association for Computational Linguistics: EACL 2023*. 910–924.
- [98] Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher. 2020. Evaluating the factual consistency of abstractive text summarization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)*. Online. Association for Computational Linguistics, 9332–9346.
- [99] Alex Kulesza and Ben Taskar. 2012. Determinantal point processes for machine learning. *Foundations and Trends in Machine Learning* 5, 2–3 (2012), 123–286.
- [100] Philippe Laban, Tobias Schnabel, Paul N. Bennett, and Marti A. Hearst. 2022. SummaC: Re-visiting NLI-based models for inconsistency detection. *Transactions of the Association for Computational Linguistics* 10 (2022), 163–177.
- [101] Philippe Laban, Wojciech Kryscinski, Divyansh Agarwal, Alexander Fabbri, Caiming Xiong, Shafiq Joty, and Chien-Sheng Wu. 2023. SummEdits: Measuring LLM ability at factual reasoning through the lens of summarization. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023)*. Singapore. Association for Computational Linguistics, 9662–9676.
- [102] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. ALBERT: A lite BERT for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- [103] Md Tahmid Rahman Laskar, Xue-Yong Fu, Cheng Chen, and Shashi Bhushan TN. 2023. Building real-world meeting summarization systems using large language models: A practical perspective. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track*. Singapore. Association for Computational Linguistics, 343–352.
- [104] Md Tahmid Rahman Laskar, Mizanur Rahman, Israt Jahan, Enamul Hoque, and Jimmy Huang. 2023. Can large language models fix data annotation errors? An empirical study using debatepedia for query-focused text summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. Singapore. Association for Computational Linguistics, 10999–11013.
- [105] Logan Lebanoff, Kaiqiang Song, and Fei Liu. 2018. Adapting the neural encoder-decoder framework from single to multi-document summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP 2018)*. Brussels, Belgium. Association for Computational Linguistics, 4131–4141.
- [106] Logan Lebanoff, Kaiqiang Song, Franck Dernoncourt, Doo Soon Kim, Seokhwan Kim, Walter Chang, and Fei Liu. 2019. Scoring sentence singletons and pairs for abstractive summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy. Association for Computational Linguistics, 2175–2189.
- [107] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*. Online. Association for Computational Linguistics, 7871–7880.
- [108] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (NeurIPS 2020)*. 9459–9474.
- [109] Raymond Li, Wen Xiao, Lanjun Wang, Gabriel Murray, and Giuseppe Carenini. 2021. Human interpretation and exploitation of self-attention patterns in transformers: A case study in extractive summarization. *arXiv preprint arXiv:2112.05364*.

- [110] Taiji Li, Zhi Li, and Yin Zhang. 2024. Improving faithfulness of large language models in summarization via sliding generation and self-consistency. *arXiv preprint arXiv:2407.21443*.
- [111] Yixiao Li, Yifan Yu, Chen Liang, Nikos Karampatziakis, Pengcheng He, Weizhu Chen, and Tuo Zhao. 2023. LoftQ: LoRA-fine-tuning-aware quantization for large language Models. *arXiv preprint arXiv:2310.08659*.
- [112] Zekai Li, Yanxia Qin, Qian Liu, and Min-Yen Kan. 2024. ISQA: Informative factuality feedback for scientific summarization. *arXiv preprint arXiv:2404.13246*.
- [113] Zizhong Li, Haopeng Zhang, and Jiawei Zhang. 2024. Unveiling the magic: Investigating attention distillation in retrieval-augmented generation. In *Proceedings of the Association for Computational Linguistics: NAACL 2024*.
- [114] Xinnian Liang, Shuangzhi Wu, Mu Li, and Zhoujun Li. 2021. Improving unsupervised extractive summarization with facet-aware modeling. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*. Online. Association for Computational Linguistics, 1685–1697.
- [115] Chin-Yew Lin and Eduard Hovy. 2003. Automatic evaluation of summaries using n-gram co-occurrence statistics. In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*. 150–157.
- [116] Hui Lin and Jeff Bilmes. 2011. A class of submodular functions for document summarization. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*. 510–520.
- [117] Jingyang Lin, Hang Hua, Ming Chen, Yikang Li, Jenhao Hsiao, Chiuman Ho, and Jiebo Luo. 2023. VideoXum: Cross-modal visual and textual summarization of videos. *IEEE Transactions on Multimedia* 26 (2023), 5548–5560.
- [118] Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. TruthfulQA: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL 2022)*. Dublin, Ireland. Association for Computational Linguistics, 3214–3252.
- [119] Fei Liu, Jeffrey Flanigan, Sam Thomson, Norman Sadeh, and Noah A. Smith. 2018. Toward abstractive summarization using semantic representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP 2018)*. Brussels, Belgium. Association for Computational Linguistics, 4095–4106.
- [120] Mengsha Liu, Daoyuan Chen, Yaliang Li, Guian Fang, and Ying Shen. 2024. ChartThinker: A contextual chain-of-thought approach to optimized chart summarization. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*. Torino, Italy, 3057–3074.
- [121] Peter J. Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. 2018. Generating Wikipedia by summarizing long sequences. In *Proceedings of the 6th International Conference on Learning Representations (ICLR 2018)*.
- [122] Shuaiqi Liu, Jiaxing Cao, Zhiyuan Deng, Wenxuan Zhao, Ruosong Yang, Zhiwen Wen, and Philip S. Yu. 2023. Neural abstractive summarization for long text and multiple tables. *IEEE Transactions on Knowledge and Data Engineering* 36, 6 (2023), 6513–6525.
- [123] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.
- [124] Ye Liu, Jianguo Zhang, Yao Wan, Congying Xia, Lifang He, and Philip S. Yu. 2021. HETFORMER: Heterogeneous transformer with sparse attention for long-text extractive summarization. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP 2021)*. Online and Punta Cana, Dominican Republic. Association for Computational Linguistics, 146–154.
- [125] Yixin Liu, Pengfei Liu, Dragomir Radev, and Graham Neubig. 2022. BRIO: Bringing order to abstractive summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL 2022)*. Dublin, Ireland. Association for Computational Linguistics, 2890–2903.
- [126] Yixin Liu, Budhaditya Deb, Milagro Teruel, Aaron Halfaker, Dragomir Radev, and Ahmed Hassan Awadallah. 2023. On improving summarization factual consistency from natural language feedback. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL 2023)*. Toronto, Canada. Association for Computational Linguistics, pages 15144–15161.
- [127] Yixin Liu, Alexander R. Fabbri, Jiawen Chen, Yilun Zhao, Simeng Han, Shafiq Joty, Pengfei Liu, Dragomir Radev, Chien-Sheng Wu, and Arman Cohan. 2024. Benchmarking generation and evaluation capabilities of large language models for instruction controllable summarization. In *Findings of the Association for Computational Linguistics: NAACL 2024*. Mexico City, Mexico. Association for Computational Linguistics, 4064–4082.
- [128] Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. G-Eval: NLG evaluation using GPT-4 with better human alignment. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023)*. Singapore. Association for Computational Linguistics, 2511–2522.
- [129] Yixin Liu, Kejian Shi, Katherine S. He, Longtian Ye, Alexander R. Fabbri, Pengfei Liu, Dragomir Radev, and Arman Cohan. 2024. On learning to summarize with large language models as references. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2024)*. Mexico City, Mexico. Association for Computational Linguistics, 8647–8664.

- [130] Yang Liu and Mirella Lapata. 2019. Hierarchical transformers for multi-document summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 5070.
- [131] Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. 3730–3740.
- [132] Yixin Liu and Pengfei Liu. 2021. SimCLS: A simple framework for contrastive learning of abstractive summarization. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*. 1065–1072.
- [133] Yao Lu, Yue Dong, and Laurent Charlin. 2020. Multi-XScience: A large-scale dataset for extreme multi-document summarization of scientific articles. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP'20)*. 8068–8074.
- [134] Hans Peter Luhn. 1958. The automatic creation of literature abstracts. *IBM Journal of Research and Development* 2, 2 (1958), 159–165.
- [135] Zheheng Luo, Qianqian Xie, and Sophia Ananiadou. 2023. ChatGPT as a factual inconsistency evaluator for abstractive text summarization. *arXiv preprint arXiv:2303.15621*.
- [136] Qing Lyu, Marianna Apidianaki, and Chris Callison-Burch. 2024. Towards faithful model explanation in NLP: A survey. *Computational Linguistics* 50, 2(2024), 657–723.
- [137] Chao Ma, Zonglin Wu, Jing Wang, Shuai Xu, Yuxuan Wei, Zhengliang Liu, Xiaowei Jiang, Liansheng Guo, Xiaowei Cai, Shaoting Zhang, and Tianming Zhang. 2023. ImpressionGPT: An iterative optimizing framework for radiology report summarization with ChatGPT. *arXiv preprint arXiv:2304.12655*.
- [138] Xinyin Ma, Gongfan Fang, and Xinchao Wang. 2023. LLM-Pruner: On the structural pruning of large language models. In *Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS 2023)*.
- [139] Inderjeet Mani and Mark T. Maybury. 1999. *Advances in Automatic Text Summarization*. MIT Press.
- [140] Qianren Mao, Jianxin Li, Hao Peng, Shizhu He, Lihong Wang, Philip S. Yu, and Zheng Wang. 2022. Fact-driven abstractive summarization by utilizing multi-granular multi-relational knowledge. *IEEE/ACM Trans. Audio, Speech and Lang. Proc.* 30 (2022), 1665–1678. <https://doi.org/10.1109/TASLP.2022.3161157>
- [141] Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On faithfulness and factuality in abstractive summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*. Online. Association for Computational Linguistics, 1906–1919.
- [142] Ryan McDonald. 2007. A study of global inference algorithms in multi-document summarization. In *Proceedings of the European Conference on Information Retrieval*. Springer, 557–564.
- [143] Rada Mihalcea and Paul Tarau. 2004. TextRank: Bringing order into text. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*. 404–411.
- [144] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- [145] Nishant Mishra, Gaurav Sahu, Iacer Calixto, Ameen Abu-Hanna, and Issam Laradji. 2023. LLM aided semi-supervision for efficient extractive dialog summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. Singapore. Association for Computational Linguistics, 10002–10009.
- [146] Muhammad F. Mridha, Aklima Akter Lima, Kamruddin Nur, Sujoy Chandra Das, and Mahmud Hasan. 2021. A survey of automatic text summarization: Progress, process and challenges. *IEEE Access* 9 (2021), 156043–156070.
- [147] Sourjit Mukherjee, Anubhav Jangra, Sriparna Saha, and Adam Jatowt. 2022. Topic-aware multimodal summarization. In *Findings of the Association for Computational Linguistics: AACL-IJCNLP 2022*. Online only. Association for Computational Linguistics, 387–398.
- [148] Nishanth Nakshatri, Siyi Liu, Sihao Chen, Dan Roth, Dan Goldwasser, and Daniel Hopkins. 2023. Using LLM for improving key event discovery: Temporal-guided news stream clustering with event summaries. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. Singapore. Association for Computational Linguistics, 4162–4173.
- [149] Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gülçehre, and Bing Xiang. 2016. Abstractive text summarization using sequence-to-sequence RNNs and beyond. In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning (CoNLL 2016)*. Berlin, Germany. Association for Computational Linguistics, 280–290.
- [150] Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2017. SummaRuNNer: A recurrent neural network based sequence model for extractive summarization of documents. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI 2017)*. San Francisco, California. AAAI Press, 3075–3091.
- [151] Feng Nan, Ramesh Nallapati, Zhiguo Wang, Cicero Nogueira dos Santos, Henghui Zhu, Dejiao Zhang, Kathleen McKeown, and Bing Xiang. 2021. Entity-level factual consistency of abstractive text summarization. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2021)*. Online. Association for Computational Linguistics, 2727–2738.

- [152] Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. Don't give me the details, just the summary! Topic-aware convolutional neural networks for extreme summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP 2018)*. Brussels, Belgium. Association for Computational Linguistics, 1797–1807.
- [153] Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. Ranking sentences for extractive summarization with reinforcement learning. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. New Orleans, Louisiana. Association for Computational Linguistics, 1747–1759.
- [154] Roberto Navigli, Simone Conia, and Björn Ross. 2023. Biases in large language models: Origins, inventory, and discussion. *ACM Journal of Data and Information Quality* 15, 2 (2023), 1–21.
- [155] Marcel Nawrath, Agnieszka Nowak, Tristan Ratz, Danilo C. Walenta, Juri Opitz, Leonardo F. R. Ribeiro, João Sedoc, et al. 2024. On the role of summary content units in text summarization evaluation. *arXiv preprint arXiv:2404.01701*.
- [156] Ani Nenkova and Rebecca J. Passonneau. 2004. Evaluating content selection in summarization: The pyramid method. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*. 145–152.
- [157] OpenAI, Greg Brockman, Sandhini Bubeck, Tom Chan, Aakanksha Chowdhery, Jacob Dell, et al. 2023. GPT-4 technical report. *arXiv:2303.08774*. Retrieved from <https://arxiv.org/abs/2303.08774>
- [158] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* 35 (2022), 27730–27744.
- [159] Paul Over, Hoa Dang, and Donna Harman. 2007. DUC in context. *Information Processing & Management* 43, 6 (2007), 1506–1520.
- [160] Romain Paulus, Caiming Xiong, and Richard Socher. 2018. A deep reinforced model for abstractive summarization. In *International Conference on Learning Representations (ICLR 2018)*.
- [161] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha, Qatar. Association for Computational Linguistics, 1532–1543.
- [162] Maxime Peyrard, Teresa Botschen, and Iryna Gurevych. 2017. Learning to score system summaries for better content selection evaluation. In *Proceedings of the Workshop on New Frontiers in Summarization*. Copenhagen, Denmark. Association for Computational Linguistics, 74–84.
- [163] Sebastian Porsdam Mann, Brian D. Earp, Nikolaj Möller, and Suren Vynn. 2023. AUTOGEN: A personalized large language model for academic enhancement-Ethics and proof of principle. *American Journal of Bioethics* 23, 10 (2023), 1327–1339.
- [164] Dongqi Pu and Vera Demberg. 2023. ChatGPT vs human-authored text: Insights into controllable text summarization and sentence style transfer. *arXiv:2306.07799*. Retrieved from <https://arxiv.org/abs/2306.07799>
- [165] Dongqi Pu and Vera Demberg. 2024. RST-LoRA: A discourse-aware low-rank adaptation for long document abstractive summarization. *arXiv:2405.00657*. Retrieved from <https://arxiv.org/abs/2405.00657>
- [166] Xiao Pu, Mingqi Gao, and Xiaojun Wan. 2023. Summarization is (almost) dead. *arXiv preprint arXiv:2309.09558*.
- [167] Weizhen Qi, Yu Yan, Yeyun Gong, Dayiheng Liu, Nan Duan, Jiusheng Chen, Ruofei Zhang, and Ming Zhou. 2020. ProphetNet: Predicting future N-gram for sequence-to-sequence pre-training. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. 2401–2410.
- [168] Leigang Qu, Meng Liu, Da Cao, Liqiang Nie, and Qi Tian. 2020. Context-aware multi-view summarization network for image-text matching. In *Proceedings of the 28th ACM International Conference on Multimedia (ACM MM 2020)*. 1047–1055.
- [169] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. *OpenAI Technical Report*.
- [170] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research* 21, 140 (2020), 1–67.
- [171] Mathieu Ravaut, Aixin Sun, Nancy F. Chen, and Shafiq Joty. 2023. On context utilization in summarization with large language models. *arXiv preprint arXiv:2310.10570*.
- [172] Mathieu Ravaut, Hailin Chen, Ruochen Zhao, Chengwei Qin, Shafiq Joty, and Nancy F. Chen. 2023. PromptSum: Parameter-efficient controllable abstractive summarization. *arXiv preprint arXiv:2308.03117*.
- [173] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. 3982–3992.
- [174] Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP 2015)*. Lisbon, Portugal. Association for Computational Linguistics, 379–389.

- [175] Marco Ferdinand Salchner and Adam Jatowt. 2022. A survey of automatic text summarization using graph neural networks. In *Proceedings of the 29th International Conference on Computational Linguistics*. 6139–6150.
- [176] Evan Sandhaus. 2008. The New York Times annotated corpus. *Linguistic Data Consortium, Philadelphia* 6, 12 (2008), e26752.
- [177] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- [178] Ritesh Sarkhel, Moniba Keymanesh, Arnab Nandi, and Srinivasan Parthasarathy. 2020. Interpretable multi-headed attention for abstractive summarization at controllable lengths. In *Proceedings of the 28th International Conference on Computational Linguistics (COLING 2020)*. Barcelona, Spain (Online). International Committee on Computational Linguistics, 6871–6882.
- [179] Thomas Scialom, Sylvain Lamprier, Benjamin Piwowarski, and Jacopo Staiano. 2019. Answers Unite! Unsupervised metrics for reinforced summarization. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China. Association for Computational Linguistics, 3246–3256.
- [180] Thomas Scialom, Paul-Alexis Dray, Patrick Gallinari, Sylvain Lamprier, Benjamin Piwowarski, Jacopo Staiano, and Alex Wang. 2021. QuestEval: Summarization asks for fact-based evaluation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP 2021)*. 6594–6604.
- [181] Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vancouver, Canada. Association for Computational Linguistics, 1073–1083.
- [182] Royal Sequiera, Luchen Tan, and Jimmy Lin. 2018. Overview of the TREC 2018 real-time summarization track. In *Proceedings of the Twenty-Seventh Text REtrieval Conference (TREC 2018)*, Gaithersburg, Maryland, USA, NIST.
- [183] Chantal Shaib, Millicent Li, Sebastian Joseph, Iain Marshall, Junyi Jessie Li, and Byron Wallace. 2023. Summarizing, simplifying, and synthesizing medical evidence using GPT-3 (with varying success). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Toronto, Canada. Association for Computational Linguistics, 1387–1407.
- [184] Ori Shapira, Hadar Ronen, Meni Adler, Yael Amsterdamer, Judit Bar-Ilan, and Ido Dagan. 2017. Interactive abstractive summarization for event news tweets. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Copenhagen, Denmark. Association for Computational Linguistics, 109–114.
- [185] Ori Shapira, Ramakanth Pasunuru, Hadar Ronen, Mohit Bansal, Yael Amsterdamer, and Ido Dagan. 2021. Extending multi-document summarization evaluation to the interactive setting. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2021)*. 657–677.
- [186] Eva Sharma, Chen Li, and Lu Wang. 2019. BIGPATENT: A large-scale dataset for abstractive and coherent summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 2204–2213.
- [187] Chenhui Shen, Liying Cheng, Xuan-Phi Nguyen, Yang You, and Lidong Bing. 2023. Large language models are not yet human-level evaluators for abstractive summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. 4215–4233.
- [188] Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: An autonomous agent with dynamic memory and self-reflection. *arXiv preprint arXiv:2303.11366*.
- [189] Kaiqiang Song, Lin Zhao, and Fei Liu. 2018. Structure-infused copy mechanisms for abstractive summarization. In *Proceedings of the 27th International Conference on Computational Linguistics (COLING 2018)*, Santa Fe, New Mexico, USA. Association for Computational Linguistics, 1717–1729.
- [190] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. MASS: Masked sequence to sequence pre-training for language generation. *arXiv preprint arXiv:1905.02450*.
- [191] Mingjie Sun, Zhuang Liu, Anna Bair, and J. Zico Kolter. 2023. A simple and effective pruning approach for large language models. *arXiv preprint arXiv:2306.11695*.
- [192] Shichao Sun, Ruifeng Yuan, Ziqiang Cao, Wenjie Li, and Pengfei Liu. 2024. Prompt chaining or stepwise prompt? Refinement in text summarization. *arXiv preprint arXiv:2404.12345*.
- [193] Shichao Sun and Wenjie Li. 2021. Alleviating exposure bias via contrastive learning for abstractive text summarization. arXiv:2108.11846. Retrieved from <https://arxiv.org/abs/2108.11846>
- [194] Weisong Sun, Chunrong Fang, Yudu You, Yun Miao, Yi Liu, Yuekang Li, Gelei Deng, Shenghan Huang, Yuchen Chen, Qunjun Zhang, Hanwei Qian, Yang Liu, and Zhenyu Chen. 2023. Automatic code summarization via ChatGPT: How far are we? *arXiv preprint arXiv:2305.12865*.
- [195] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems (NeurIPS 2014)*. 3104–3112.

- [196] Hiroya Takamura and Manabu Okumura. 2009. Text summarization model based on maximum coverage problem and its variant. In *Proceedings of the 12th Conference of the European Chapter of the ACL (EACL '09)*. 781–789.
- [197] Derek Tam, Anisha Mascarenhas, Shiyue Zhang, Sarah Kwan, Mohit Bansal, and Colin Raffel. 2023. Evaluating the factual consistency of large language models through news summarization. In *Findings of the Association for Computational Linguistics: ACL 2023*. 4127–4144.
- [198] Liyan Tang, Igor Shalyminov, Amy Wing-mei Wong, Jon Burnsky, Jake W. Vincent, Yu'an Yang, Siffi Singh, et al. 2024. TofuEval: Evaluating hallucinations of LLMs on topic-focused dialogue summarization. *arXiv preprint arXiv:2402.13249*.
- [199] Liyan Tang, Zhaoyi Sun, Betina R. S. Idnay, Jordan G. Nestor, Amani Soroush, Pablo Adolfo Elias, Z. P. Xu, Y. Ding, Greg Durrett, Justin F. Rousseau, Cindy Weng, and Yalei Peng. 2023. Evaluating large language models on medical evidence summarization. *NPJ Digital Medicine* 6, Article 134 (2023).
- [200] Yi Tay, Mostafa Dehghani, Viet Q. Tran, Xavier Garcia, Jason Wei, Xuezhi Wang, Hyung Won Chung, Aakanksha Chowdhery, Sharan Narang, Adam Roberts, Denny Zhou, and Donald Metzler. 2022. UL2: Unifying language learning paradigms. *arXiv preprint arXiv:2205.05131*.
- [201] Gemini Team et al. 2023. Gemini: A family of highly capable multimodal models. arXiv:2312.11805. Retrieved from <https://arxiv.org/abs/2312.11805>
- [202] Arun James Thirunavukarasu, Daniel Shu Wei Ting, Karthikeyan Elangovan, Luis Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. 2023. Large language models in medicine. *Nature Medicine* 29, 8 (2023), 1930–1940.
- [203] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, et al. 2023. LLaMA: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- [204] Dave Van Veen et al. 2023. Clinical text summarization: Adapting large language models can outperform human experts. *Research Square* (2023).
- [205] Oleg Vasilyev, Vedant Dharnidharka, and John Bohannon. 2020. Fill in the BLANC: Human-free quality estimation of document summaries. In *Proceedings of the First Workshop on Evaluation and Comparison of NLP Systems (Eval4NLP 2020)*. Online. Association for Computational Linguistics, 11–20.
- [206] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS 2017)*. 5998–6008.
- [207] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *Proceedings of the 6th International Conference on Learning Representations (ICLR 2018)*.
- [208] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Advances in Neural Information Processing Systems (NeurIPS 2015)*. 2692–2700.
- [209] Yao Wan, Zhou Zhao, Min Yang, Guandong Xu, Hao Ying, Jian Wu, and Philip S. Yu. 2018. Improving automatic source code summarization via deep reinforcement learning. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering (ASE 2018)*. 397–407.
- [210] Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020. Asking and answering questions to evaluate the factual consistency of summaries. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*. 7516–7531.
- [211] Danqing Wang, Pengfei Liu, Yining Zheng, Xipeng Qiu, and Xuanjing Huang. 2020. Heterogeneous graph neural networks for extractive document summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*. 6209–6219.
- [212] Fei Wang, Kaiqiang Song, Hongming Zhang, Lifeng Jin, Sangwoo Cho, Wenlin Yao, Xiaoyang Wang, Muhao Chen, and Dong Yu. 2022. Salience allocation as guidance for abstractive summarization. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022)*. Abu Dhabi, United Arab Emirates. Association for Computational Linguistics, 6094–6106.
- [213] Jiaan Wang, Yunlong Liang, Fandong Meng, Beiqi Zou, Zhixu Li, Jianfeng Qu, and Jie Zhou. 2023. Zero-shot cross-lingual summarization via large language models. In *Proceedings of the 4th New Frontiers in Summarization Workshop*. 12–23.
- [214] Wenhua Wang, Yuqun Zhang, Yulei Sui, Yao Wan, Zhou Zhao, Jian Wu, Philip S. Yu, and Guandong Xu. 2020. Reinforcement-learning-guided source code summarization using hierarchical attention. *IEEE Transactions on Software Engineering* 48, 11 (2020), 4139–4156.
- [215] Yiming Wang, Zhuosheng Zhang, and Rui Wang. 2023. Element-aware summarization with large language models: Expert-aligned evaluation and chain-of-thought method. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL 2023)*. 8640–8665.
- [216] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems (NeurIPS 2022)*. 1–21.

- [217] Chenxi Whitehouse, Fantine Huot, Jasmijn Bastings, Mostafa Dehghani, Chia-chen Lin, and Mirella Lapata. 2023. Parameter-efficient multilingual summarisation: An empirical study. In *Findings of the North American Chapter of the Association for Computational Linguistics: NAACL 2024*. 1202–1228.
- [218] Aji Prasetya Wibawa and Fachrul Kurniawan. 2024. A survey of text summarization: Techniques, evaluation and challenges. *Natural Language Processing Journal* 7 (2024), 100070.
- [219] Jimmy Wu, Rika Antonova, Adam Kan, Marion Lepert, Andy Zeng, Shuran Song, Jeannette Bohg, Szymon Rusinkiewicz, and Thomas Funkhouser. 2023. TidyBot: Personalized robot assistance with large language models. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2023)*. 3546–3553.
- [220] Ning Wu, Ming Gong, Linjun Shou, Shining Liang, and Daxin Jiang. 2023. Large language models are diverse role-players for summarization evaluation. In *Proceedings of the 12th CCF International Conference on Natural Language Processing and Chinese Computing (NLPCC 2023)*. 637–649.
- [221] Yunshu Wu, Hayate Iso, Pouya Pezeshkpour, Nikita Bhutani, and Estevam Hruschka. 2023. Less is more for long document summary evaluation by LLMs. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2024), Volume 2: Short Papers*. St. Julian's, Malta. Association for Computational Linguistics, 330–343.
- [222] Yu Xia, Xu Liu, Tong Yu, Sungchul Kim, Ryan A. Rossi, Anup Rao, Tung Mai, and Shuai Li. 2024. Hallucination diversity-aware active learning for text summarization. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2024)*. 14620–14635.
- [223] Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. 2023. SmoothQuant: Accurate and efficient post-training quantization for large language models. In *Proceedings of the 40th International Conference on Machine Learning (ICML 2023)*. 14420–14442.
- [224] Min Xiao, Junnan Zhu, Haitao Lin, Yu Zhou, and Chengqing Zong. 2023. CFSum: A coarse-to-fine contribution network for multimodal summarization. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Toronto, Canada. Association for Computational Linguistics, 8538–8553.
- [225] Wen Xiao, Iz Beltagy, Giuseppe Carenini, and Arman Cohan. 2022. PRIMERA: Pyramid-based masked sentence pre-training for multi-document summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL 2022)*. Dublin, Ireland. Association for Computational Linguistics, 5245–5263.
- [226] Wen Xiao and Giuseppe Carenini. 2019. Extractive summarization of long documents by combining global and local context. arXiv:1909.08089. Retrieved from <https://arxiv.org/abs/1909.08089>
- [227] Wen Xiao and Giuseppe Carenini. 2020. Systematically exploring redundancy reduction in summarizing long documents. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*. 516–528.
- [228] Wen Xiao and Giuseppe Carenini. 2022. Entity-based spancopy for abstractive summarization to improve the factual consistency. arXiv:2209.03479. Retrieved from <https://arxiv.org/abs/2209.03479>
- [229] Frank F. Xu, Uri Alon, Graham Neubig, and Vincent J. Hellendoorn. 2022. A systematic evaluation of large language models of code. In *Proceedings of the 6th Annual Symposium on Machine Programming (MAPS 2022)*. 1–12.
- [230] Hongyan Xu, Hongtao Liu, Zhepeng Lv, Qing Yang, and Wenjun Wang. 2023. Sentiment-aware review summarization with personalized multi-task fine-tuning. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM 2023)*. 2826–2835.
- [231] Jiacheng Xu, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Discourse-aware neural extractive text summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*. Online. Association for Computational Linguistics, 5021–5031.
- [232] Liyan Xu, Zhenlin Su, Mo Yu, Jin Xu, Jinho D. Choi, Jie Zhou, and Fei Liu. 2024. Identifying factual inconsistency in summaries: Towards effective utilization of large language model. In *Findings of the Association for Computational Linguistics: EMNLP 2024*. Miami, Florida, USA. Association for Computational Linguistics, 14626–14641.
- [233] Shusheng Xu, Xingxing Zhang, Yi Wu, and Furu Wei. 2022. Sequence level contrastive learning for text summarization. In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2022)*. 11556–11565.
- [234] Xiaohan Xu, Ming Li, Chongyang Tao, Tao Shen, Reynold Cheng, Jinyang Li, Can Xu, Dacheng Tao, and Tianyi Zhou. 2024. A survey on knowledge distillation of large language models. *arXiv preprint arXiv:2402.13116*.
- [235] Yichong Xu, Ruochen Xu, Dan Iter, Yang Liu, Shuohang Wang, Chenguang Zhu, and Michael Zeng. 2023. InheritSumm: A general, versatile and compact summarizer by distilling from GPT. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. Singapore. Association for Computational Linguistics, 13879–13894.
- [236] Shweta Yadav, Deepak Gupta, Asma Ben Abacha, and Dina Demner-Fushman. 2021. Reinforcement learning for abstractive question summarization with question-aware semantic rewards. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. Online. Association for Computational Linguistics, 249–255.

- [237] Min Yang, Qiang Qu, Ying Shen, Qiao Liu, Wei Zhao, and Jia Zhu. 2018. Aspect and sentiment aware abstractive review summarization. In *Proceedings of the 27th International Conference on Computational Linguistics (COLING 2018)*. Santa Fe, New Mexico, USA, 1110–1120.
- [238] Saellyne Yang, Jisu Yim, Juho Kim, and Hijung Valentina Shin. 2022. CatchLive: Real-time summarization of live streams with stream content and interaction data. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI'22), April 29 May 5, 2022, New Orleans, LA, USA*. Association for Computing Machinery, Article 147, 1–20.
- [239] Xiaoxuan Yang, Bonan Yan, Hai Li, and Yiran Chen. 2020. ReTransformer: ReRAM-based processing-in-memory architecture for transformer acceleration. In *Proceedings of the 39th International Conference on Computer-Aided Design (ICCAD 2020)*. Article 147 (2020), 19. DOI: [10.1145/3400302.3415640](https://doi.org/10.1145/3400302.3415640)
- [240] Xiaoxuan Yang, Huanrui Yang, Jingchi Zhang, Hai Helen Li, and Yiran Chen. 2022. On building efficient and robust neural network designs. In *Proceedings of the 56th Asilomar Conference on Signals, Systems, and Computers (ASILOMAR 2022)*. 317–321.
- [241] Xianjun Yang, Yan Li, Xinlu Zhang, Haifeng Chen, and Wei Cheng. 2023. Exploring the limits of ChatGPT for query or aspect-based text summarization. *arXiv preprint arXiv:2302.08081*.
- [242] Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. Graph convolutional networks for text classification. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI 2019)*. 7370–7377.
- [243] Zonghai Yao, Benjamin J. Schloss, and Sai P. Selvaraj. 2023. Improving summarization with human edits. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023)*. Singapore. Association for Computational Linguistics, 2604–2620.
- [244] Michihiro Yasunaga, Rui Zhang, Kshitijh Meelu, Ayush Pareek, Krishnan Srinivasan, and Dragomir Radev. 2017. Graph-based neural multi-document summarization. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*. Vancouver, Canada. Association for Computational Linguistics, 452–462.
- [245] Wenpeng Yin and Yulong Pei. 2015. Optimizing sentence modeling and selection for document summarization. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*.
- [246] Wenhao Yu, Zhihan Zhang, Zhenwen Liang, Meng Jiang, and Ashish Sabharwal. 2023. Improving language models via plug-and-play retrieval feedback. *arXiv preprint arXiv:2305.14002*.
- [247] Haoan Yuan and Haopeng Zhang. 2024. DomainSum: A hierarchical benchmark for fine-grained domain shift in abstractive text summarization. arXiv:2410.15687. Retrieved from <https://arxiv.org/abs/2410.15687>
- [248] Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. BARTScore: Evaluating generated text as text generation. In *Advances in Neural Information Processing Systems (NeurIPS 2021)* 34 (2021), 27263–27277.
- [249] Haopeng Zhang, Semih Yavuz, Wojciech Kryściński, Kazuma Hashimoto, and Yingbo Zhou. 2022. Improving the faithfulness of abstractive summarization via entity coverage control. In *Findings of the Association for Computational Linguistics: NAACL 2022*. Seattle, United States. Association for Computational Linguistics, 528–535.
- [250] Haopeng Zhang, Sangwoo Cho, Kaiqiang Song, Xiaoyang Wang, Hongwei Wang, and Jiawei Zhang. 2023. Unsupervised multi-document summarization with holistic inference. In *Findings of the Association for Computational Linguistics: IJCNLP 2023*. 123–131.
- [251] Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2022. HEGEL: Hypergraph transformer for long document summarization. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. 10167–10176.
- [252] Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023. Contrastive hierarchical discourse graph for scientific document summarization. In *Proceedings of the 4th Workshop on Computational Approaches to Discourse (CODI'23)*. 37–47.
- [253] Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023. DiffuSum: Generation enhanced extractive summarization with diffusion. In *Findings of the Association for Computational Linguistics: ACL 2023*. 13089–13100.
- [254] Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023. Extractive summarization via ChatGPT for faithful summary generation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. 3270–3278.
- [255] Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023. SummIt: Iterative Text Summarization via ChatGPT. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. Singapore, 10644–10657.
- [256] Haopeng Zhang and Jiawei Zhang. 2020. Text graph transformer for document classification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)*. Association for Computational Linguistics, 8322–8332.
- [257] Jiawei Zhang et al. 2020. Graph-Bert: Only attention is needed for learning graph representations. arXiv:2001.05140. Retrieved from <https://arxiv.org/abs/2001.05140>
- [258] Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2020. PEGASUS: Pre-training with Extracted Gap-sentences for abstractive summarization. In *Proceedings of the 37th International Conference on Machine Learning*. 11328–11339.
- [259] Rui Zhang and Joel Tetreault. 2019. This email could save your life: Introducing the task of email subject line generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 446–456.

- [260] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. BERTScore: Evaluating text generation with BERT. In *Proceedings of the International Conference on Learning Representations (ICLR 2020)*.
- [261] Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B. Hashimoto. 2024. Benchmarking large language models for news summarization. In *Transactions of the Association for Computational Linguistics* 12 (2024), 33–50.
- [262] Xingxing Zhang, Mirella Lapata, Furu Wei, and Ming Zhou. 2018. Neural latent extractive document summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels, Belgium. Association for Computational Linguistics, 779–784.
- [263] Xingxing Zhang, Furu Wei, and Ming Zhou. 2019. HIBERT: Document level pre-training of hierarchical bidirectional transformers for document summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy. Association for Computational Linguistics, 5059–5069.
- [264] Yusen Zhang, Nan Zhang, Yixin Liu, Alexander Richard Fabbri, Junru Liu, Ryo Kamoi, Xiaoxin Lu, et al. 2024. Fair abstractive summarization of diverse perspectives. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*. 3404–3426.
- [265] Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lema Liu, Tingchen Fu, Xinting Huang, et al. 2023. Siren’s song in the AI Ocean: A survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219*.
- [266] Bin Zhao and Eric P. Xing. 2014. Quasi real-time summarization for consumer videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2513–2520.
- [267] Chenlong Zhao, Xiwen Zhou, Xiaopeng Xie, and Yong Zhang. 2024. Hierarchical attention graph for scientific document summarization in global and local level. In *Findings of the Association for Computational Linguistics: NAACL 2024*. 714–726.
- [268] Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M. Meyer, and Steffen Eger. 2019. MoverScore: Text generation evaluating with contextualized embeddings and earth mover distance. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China, 563–578.
- [269] Hao Zheng and Mirella Lapata. 2019. Sentence centrality revisited for unsupervised summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 6236–6247.
- [270] Ming Zhong, Danqing Wang, Pengfei Liu, Xipeng Qiu, and Xuan-Jing Huang. 2019. A closer look at data bias in neural extractive summarization models. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization (NewSum’19)*. 80–89.
- [271] Ming Zhong, Pengfei Liu, Yiran Chen, Danqing Wang, Xipeng Qiu, and Xuan-Jing Huang. 2020. Extractive summarization as text matching. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 6197–6208.
- [272] Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan, et al. 2021. QMSum: A new benchmark for query-based multi-domain meeting summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT’21)*. 5905–5921.
- [273] Ming Zhong, Yang Liu, Da Yin, Yuning Mao, Yizhu Jiao, Pengfei Liu, Chenguang Zhu, Heng Ji, and Jiawei Han. 2022. Towards a unified multi-dimensional evaluator for text generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. 2023–2038.
- [274] Ming Zhong, Yang Liu, Suyu Ge, Yuning Mao, Yizhu Jiao, Xingxing Zhang, Yichong Xu, Chenguang Zhu, Michael Zeng, and Jiawei Han. 2022. Unsupervised multi-granularity Summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2022*. 4980–4995.
- [275] Qingyu Zhou, Nan Yang, Furu Wei, Shaohan Huang, Ming Zhou, and Tiejun Zhao. 2018. Neural document summarization by jointly learning to score and select sentences. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) (ACL’18)*. 654–663.
- [276] Weixiao Zhou, Gengyao Li, Xianfu Cheng, Xinnian Liang, Junnan Zhu, Feifei Zhai, and Zhoujun Li. 2023. Multi-stage pre-training enhanced by ChatGPT for multi-scenario multi-domain dialogue summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. 6893–6908.
- [277] Chenguang Zhu, William Hinthorn, Ruochen Xu, Qingkai Zeng, Michael Zeng, Xuedong Huang, and Meng Jiang. 2021. Enhancing factual consistency of abstractive summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 718–733.
- [278] Chenguang Zhu, Yang Liu, Jie Mei, and Michael Zeng. 2021. MediaSum: A large-scale media interview dataset for dialogue summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 5927–5934.

- [279] Junnan Zhu, Haoran Li, Tianshang Liu, Yu Zhou, Jiajun Zhang, and Chengqing Zong. 2018. MSMO: Multimodal summarization with multimodal output. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 4154–4164.
- [280] Junnan Zhu, Yu Zhou, Jiajun Zhang, Haoran Li, Chengqing Zong, and Changliang Li. 2020. Multimodal summarization with guidance of multimodal reference. In *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 05 (2020), 9749–9756.

Received 20 June 2024; revised 12 February 2025; accepted 1 April 2025