

# Low-Resource Cross-Lingual Summarization through Few-Shot Learning with Large Language Models

Gyutae Park, Seojin Hwang, Hwanhee Lee<sup>†</sup>

Department of Artificial Intelligence, Chung-Ang University, Seoul, Korea  
{pkt0401, swiftie1230, hwanheelee}@cau.ac.kr

## Abstract

Cross-lingual summarization (XLS) aims to generate a summary in a target language different from the source language document. While large language models (LLMs) have shown promising zero-shot XLS performance, their few-shot capabilities on this task remain unexplored, especially for low-resource languages with limited parallel data. In this paper, we investigate the few-shot XLS performance of various models, including Mistral-7B-Instruct-v0.2, GPT-3.5, and GPT-4. Our experiments demonstrate that few-shot learning significantly improves the XLS performance of LLMs, particularly GPT-3.5 and GPT-4, in low-resource settings. However, the open-source model Mistral-7B-Instruct-v0.2 struggles to adapt effectively to the XLS task with limited examples. Our findings highlight the potential of few-shot learning for improving XLS performance and the need for further research in designing LLM architectures and pre-training objectives tailored for this task. We provide a future work direction to explore more effective few-shot learning strategies and to investigate the transfer learning capabilities of LLMs for cross-lingual summarization.

## 1 Introduction

Cross-Lingual Summarization (XLS) is a task that involves generating a summary in a target language different from the source document’s language. It is a complex natural language processing task that requires performing both summarization and machine translation simultaneously. This task is much more challenging than single-language summarization, as it involves overcoming the differences between languages while effectively extracting and compressing key information. Generally, there are two types of pipelines for XLS systems (Leuski et al., 2003; Orăsan and Chiorean, 2008): *summarize-then-translate* and




Figure 1: An example of cross-lingual summarization.

*translate-then-summarize*. The former summarizes the source text first and then translates it, while the latter translates the source text first and then summarizes it.

However, these pipelines have the disadvantage that errors occurring at each stage can propagate and accumulate, potentially degrading the final performance. To resolve the issues of these pipelines, there has been a lot of research on the end-to-end approach (Zhu et al., 2019; Bai et al., 2021), known as the direct method, which generates the target language summary directly from the source document in a single step. The direct method can mitigate the error propagation problem compared to traditional pipelines and enable more efficient learning.

In parallel with the development of end-to-end approaches, Large Language Models (LLMs) such as GPT-3.5 and GPT-4 (OpenAI, 2023) have demonstrated strong performance in various natural language processing tasks. It is known that these models can significantly improve performance through few-shot learning with only a small number of examples. (Brown et al., 2020) However, directly applying these models to the XLS task is still challenging. Particularly for low-resource languages, where it is difficult to build large-scale

<sup>†</sup>Corresponding author.

parallel corpora, the lack of data makes it difficult to fully utilize the performance of pre-trained language models, acting as a major obstacle in XLS research. (Ladhak et al., 2020) Hence, this study aims to explore the XLS performance of various LLMs using a few-shot learning approach through in-context learning, focusing on the direct method.

Our findings demonstrate that few-shot learning enables LLMs, particularly GPT-3.5 and GPT-4, to achieve competitive performance in cross-lingual summarization tasks for low-resource language settings. The results also highlight open-source models’ challenges in adapting to the XLS task with limited parallel data. These findings emphasize the potential of few-shot learning in enhancing cross-lingual summarization capabilities and the need for further research in developing effective few-shot strategies and architectures for low-resource languages.

## 2 Related Work

Cross-lingual summarization (XLS) has undergone significant evolution, shifting from early pipeline approaches like *summarize-then-translate* (Orăsan and Chiorean, 2008; Wan et al., 2010) and *translate-then-summarize* (Leuski et al., 2003; Boudin et al., 2011) to more sophisticated methods utilizing multilingual pre-trained models. These pipelines were initially dominant due to their simplicity but were plagued by error propagation and the limitations inherent to sequential processing tasks. (Parnell et al., 2024) The advent of multilingual pre-trained models such as mBART (Lewis et al., 2020) and mT5 (Xue et al., 2021) marked a transformative shift towards end-to-end approaches, directly generating summaries in the target language and substantially mitigating error propagation issues.

In recent years, the emergence of large language models (LLMs) has revolutionized the field of natural language processing, including XLS. Especially, widely used LLMs like GPT-3.5 and GPT-4 have demonstrated remarkable zero-shot learning capabilities across various tasks. (Brown et al., 2020; Qin et al., 2023; Bubeck et al., 2023) However, the exploration of LLMs in the context of XLS is still in its early stages, with limited research on their zero-shot learning capabilities and even fewer studies focusing on their few-shot learning potential. (Wang et al., 2023) Recent studies have shown promising results in using LLMs for various NLP

tasks. (Bang et al., 2023; Yang et al., 2023; Patil and Gudivada, 2024)

However, the specific exploration of these models in XLS scenarios, particularly in the few-shot setting, remains largely unexplored. While some studies have investigated the zero-shot XLS performance of LLMs (Wang et al., 2023), there is a notable lack of research on the few-shot learning capabilities of models such as GPT-3.5, GPT-4, and multilingual LLMs in the XLS domain. Moreover, the disparity in performance between proprietary models like GPT-4 and open-source alternatives in zero-shot settings underscores the necessity for further investigation into the few-shot capabilities of LLMs. This is particularly critical to ensure that advancements in XLS are equitable and accessible across various linguistic and resource settings.

In this paper, we aim to bridge this gap by exploring the few-shot learning capabilities of LLMs in the context of XLS. We focus on the direct method, leveraging the nuanced capabilities of LLMs like GPT-3.5, GPT-4, and open-source models such as Mistral-7B-Instruct-v0.2. The Mistral-7B-Instruct-v0.2 is a 7.3B parameter model that outperforms Llama-2 13B (Touvron et al., 2023b) across all benchmarks and even surpasses Llama-1 34B (Touvron et al., 2023a) on many tasks and particularly noted for its ability to process up to 32k tokens, significantly enhancing its capability for few-shot learning by providing richer context management. (La Plateforme; Jiang et al., 2024) This open-source model has demonstrated strong performance in various natural language processing tasks and offers robust multilingual support. Our aim is to provide comprehensive insights into the practical applications and limitations of these models in low-resource languages, setting the stage for future advancements in the field.

## 3 Methods

The main objective of our research is to compare and analyze the performance of pre-trained mT5 and few-shot prompt-based GPT-3.5 and GPT-4. Then, we aim to experimentally confirm the impact of their few-shot learning on low-resource language XLS tasks. Additionally, we conduct comprehensive comparison with one of the open-source multilingual LLMs, such as the Mistral-7B-Instruct-v0.2 (La Plateforme; Jiang et al., 2024), to provide a broader perspective on the performance of different LLMs in the XLS task to provide a broader

---

### Input Prompt

---

Please summarize the following text in English

#### Example 1

**Text:** "แผนที่ของนาซาซึ่งชี้ให้เห็นว่า ระดับของไนโตรเจนไดออกไซด์มลพิษที่เกิดจากโรงงานอุตสาหกรรม มีปริมาณลดลงด้วยเช่นกันในปีนี้ สอดคล้องกับจำนวนตัวเลขของโรงงานในจีนที่ถูกลบปิด โรงงานยุติการผลิตที่ถักกันไม่ให้ไวรัสแพร่ระบาด ตัวเลขผู้ติดเชื้อไวรัสโคโรนาสายพันธุ์ใหม่ในจีนขณะนี้อยู่ที่เกือบ 80,000 ราย นับแต่เกิดการระบาดเมื่อปลายปีที่แล้ว ขณะที่พบผู้ติดเชื้อไวรัสชนิดนี้ในอีกกว่า..."  
(Nasa's map also shows that the level of nitrogen dioxide pollution caused by industrial plants has also decreased this year, in line with the number of plants in China that have been...)

**Translated summary:** "Satellite images have shown a dramatic decline in pollution levels over China, which is "at least partly" due to an economic slowdown prompted by the coronavirus, US space agency Nasa says."

#### Example 2

**Text:** "นิตยสารเพลย์บอยลงโฆษณาฉบับเดือนมีนาคม-เมษายนของตงทางทวิตเตอร์ 2017 นิตยสารเพลย์บอยซึ่งยุติการตีพิมพ์รูปเปลือยของสาวเพื่อนเล่น "เพลย์เมท" ซึ่งเป็นเอกลักษณ์ประจำตัวที่โด่งดังไปเมื่อเดือนมีนาคมปี 2016 ประกาศเปลี่ยนนโยบายแบบกลับหลังหันอีกครั้ง โดยเพลย์บอยฉบับเดือนมีนาคม-เมษายน 2017 จะลงรูปเปลือยของสาวเพลย์เมทหน้าปก โดยมีคำโฆษณา..."  
(Playboy magazine advertised its March-April issue on Twitter 2017. Playboy magazine, which stopped publishing nude photos of playmate girlfriend "Playmate", a well-known identity in...)

**Translated summary:** "Playboy magazine has announced it is bringing back nudity, reversing a decision made last year."

#### Test Text

**Text:** "สื่อรัสเซียรายงานว่า ศ.โอเลก โซโคลอฟ 63 ปี อยู่ในสภาพมีเมาและพลัดตกลงไปในแม่น้ำพร้อมกับกระเป๋าสตางค์ในแม่น้ำส่วนแขนทั้ง 2 ข้างของ น.ส.อานาสตาเซีย เยเชเชนโก 24 ปี ในเวลาต่อมา เจ้าหน้าที่ตำรวจได้เข้าตรวจสอบบ้านพักของ ศ.โซโคลอฟ ในนครเซนต์ปีเตอส์เบิร์ก และพบร่างที่ถูกหั่นแยกชิ้นส่วนของ น.ส.เยเชเชนโก เจ้าหน้าที่ตำรวจสอบสวนมรณกรรม..."  
(Russian media reported that Prof.Oleg Sokolov, 63, was intoxicated and fell into the river with a backpack containing parts of Ms.Anastasia Yeshchenko, 24, later a police officer examined...)

**Translated summary:**

---

Figure 2: Two-shot prompt construction for cross-lingual summarization from Thai to English.

perspective on the performance of different LLMs in the XLS task and assess the effectiveness of few-shot learning in mitigating the challenges posed by low-resource settings.

### 3.1 Direct Cross-Lingual Summarization

We focus on the direct cross-lingual summarization method, which generates target language summaries directly from source language documents in an end-to-end manner. Unlike traditional pipelines that involve separate summarization and translation steps, the direct approach combines these tasks into a single, unified process. This allows for a more seamless transfer of information between languages and reduces the potential for error propagation.

### 3.2 Models

We compare the performance of fine-tuned mT5, GPT-3.5, GPT-4, and Mistral-7B-Instruct-v0.2 on cross-lingual summarization tasks. For GPT-3.5 and GPT-4, specifically GPT-3.5-turbo-0125 and GPT-4-0125-preview models, and Mistral-7B-Instruct-v0.2, we evaluate their performance in zero-shot, one-shot, and two-shot settings. The Mistral-7B-Instruct-v0.2 is particularly noted for its ability to process up to 32k tokens, significantly enhancing its capability for few-shot learning by providing richer context management. This model has been shown to outperform other models like Llama-2-13B across all benchmarks, with robust multilingual support enhancing its utility for diverse linguistic datasets.

### 3.3 Few-Shot Prompt Construction

For the few-shot learning approach, we construct prompts that include several examples from the validation set. These examples are carefully selected based on their token count, ensuring that the shortest examples are used as the first and second examples in the prompt. This structured approach facilitates effective few-shot learning, even when computational resources are limited.

The direct prompts for few-shot learning, as depicted in Figure 2, are structured to provide the model with two examples of increasing complexity. Each prompt consists of two example texts and their corresponding summaries in the target language. The test text is then appended to the prompt, and the model is expected to generate a summary in English. By including these meticulously chosen examples in the prompt, we aim to provide the model with sufficient context to perform few-shot cross-lingual summarization effectively. This method allows us to explore the capabilities of large language models in low-resource settings where the availability of parallel data is limited.

## 4 Experiments

### 4.1 Datasets

We utilize the CrossSum (Bhattacharjee et al., 2023) dataset, a multilingual corpus of summaries in 45 languages. Following the definitions in (Li et al., 2023), we focus on low-resource language pairs with fewer than 1,000 parallel data points. Additionally, we include experiments with Pashto, a medium-resource language with 1,212 parallel data points, to more broadly assess the effective-

Models	Language Pair																		
	Th-En			Gu-En			Mr-En			Pa-En			Bu-En			Si-En			
	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	
mT5-Base (fine-tuning)	29.16	9.79	22.63	24.28	6.40	19.12	25.87	7.28	20.14	30.86	10.27	24.78	28.44	7.50	22.07	28.94	8.17	22.53	
GPT-3.5	Zero-shot	14.69	3.84	10.32	12.69	2.80	9.20	13.83	3.04	9.96	14.98	3.78	10.94	13.2	2.21	10.19	13.08	2.11	9.91
	One	15.18	3.95	10.99	13.0	3.12	9.25	14.91	3.73	10.79	14.62	3.93	10.80	15.80	2.43	11.91	13.39	2.34	10.11
	Two	16.56	4.59	11.6	<b>15.33</b>	<b>3.57</b>	<b>10.57</b>	<b>16.52</b>	<b>4.06</b>	<b>11.96</b>	<b>15.52</b>	3.96	<b>11.06</b>	16.48	2.83	<b>12.39</b>	14.03	2.47	9.81
GPT-4	Zero-shot	12.22	3.39	8.59	10.96	2.71	7.63	10.97	2.84	7.90	12.10	3.58	8.55	14.73	3.77	9.84	13.92	3.86	9.45
	One	14.86	4.12	10.51	13.88	3.50	9.27	14.11	3.52	10.09	13.33	4.04	9.43	16.88	<b>4.48</b>	11.41	15.03	<b>4.35</b>	10.27
	Two	<b>17.67</b>	<b>5.19</b>	<b>12.67</b>	13.63	3.49	9.27	14.97	3.73	10.38	13.65	<b>4.15</b>	9.49	<b>17.38</b>	4.32	11.73	<b>15.10</b>	4.15	<b>10.39</b>
Mistral-7B-instruct-v0.2	Zero-shot	8.91	2.12	6.42	6.15	0.68	4.76	7.63	1.55	5.76	7.65	1.62	5.99	7.59	1.06	5.68	7.23	0.99	5.40
	One	10.28	2.51	7.58	7.27	0.79	5.46	8.08	1.42	6.06	7.41	1.21	5.82	8.35	0.91	6.37	8.43	1.06	6.08
	Two	10.10	2.37	7.30	6.31	0.71	4.81	6.66	0.57	5.35	8.96	1.71	6.77	9.55	1.07	7.44	8.21	0.49	6.58

Table 1: Performance comparison of model performance metrics across various language pairs, including R1, R2, and Rouge-L scores. The language pairs are abbreviated as follows: Th-En (Thai to English), Gu-En (Gujarati to English), Mr-En (Marathi to English), Pa-En (Pashto to English), Bu-En (Burmese to English), Si-En (Sinhala to English).

Models	Language Pair																		
	En-Th			En-Gu			En-Mr			En-Pa			En-Bu			En-Si			
	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	
mT5-Base (fine-tuning)	4.59	0.94	4.40	10.14	1.19	9.37	9.38	1.22	8.94	20.12	4.62	17.55	4.85	0.55	4.72	5.62	0.39	5.13	
GPT-3.5	Zero-shot	4.78	1.48	4.53	1.40	0.0	1.40	2.05	0.0	2.05	0.0	0.0	0.0	1.70	0.26	1.70	1.22	0.0	1.22
	One	4.5	0.82	4.45	1.60	0.0	1.60	1.02	0.13	1.02	0.0	0.0	0.0	0.58	0.0	1.55	1.72	1.72	1.72
	Two	5.29	1.46	5.20	0.55	0.0	0.55	0.88	0.09	0.88	0.0	0.0	0.0	1.55	0.0	1.55	1.63	<b>1.07</b>	1.63
GPT-4	Zero-shot	6.09	1.42	5.84	1.50	0.0	1.50	1.35	0.0	1.35	0.0	0.0	0.0	2.61	0.90	2.61	3.34	<b>1.72</b>	3.34
	One	<b>9.38</b>	<b>2.25</b>	<b>9.38</b>	1.22	0.0	1.22	1.48	0.1	1.48	0.0	0.0	0.0	4.62	0.17	4.62	<b>4.36</b>	1.15	<b>4.36</b>
	Two	8.64	2.16	8.39	<b>2.05</b>	0.0	<b>2.05</b>	1.31	0.0	1.31	0.0	0.0	0.0	2.02	0.22	2.02	2.53	1.15	2.53
Mistral-7B-instruct-v0.2	Zero-shot	4.33	1.18	4.33	0.23	0.0	0.23	<b>2.37</b>	<b>0.20</b>	<b>2.37</b>	0.0	0.0	0.0	2.64	1.21	2.64	0.06	0.0	0.06
	One	4.39	1.33	4.30	0.51	0.0	0.51	1.64	0.08	1.64	0.0	0.0	0.0	<b>7.78</b>	<b>1.55</b>	<b>7.78</b>	0.74	0.0	0.74
	Two	3.15	0.38	2.94	0.51	0.0	0.51	0.55	0.0	0.55	0.0	0.0	0.0	2.43	0.01	2.10	0.42	0.0	0.42

Table 2: Performance comparison of model performance metrics across various language pairs, including R1, R2, and Rouge-L scores. The language pairs are abbreviated as follows: En-Th (English to Thai), En-Gu (English to Gujarati), En-Mr (English to Marathi), En-Pa (English to Pashto), En-Bu (English to Burmese), and En-Si (English to Sinhala).

ness of our proposed method in diverse linguistic settings. Figure 3 illustrates the distribution of the dataset across different languages for both many-to-one and one-to-many scenarios. The number of parallel data points for each language pair remains consistent in both settings. This symmetry allows us to represent the dataset distribution in a single figure, simplifying the visual representation of the data.

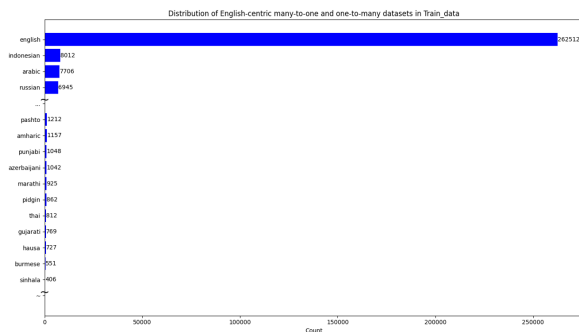


Figure 3: Distribution of English-centric many-to-one and one-to-many datasets in Train data

## 4.2 Performance Metrics

To evaluate the quality of the generated summaries, we used the ROUGE (Lin, 2004), reporting ROUGE-1/2/L (R-1,2,L). These metrics measure the overlap of unigrams, bigrams, and the longest common subsequences between the generated and reference summaries, respectively.

## 4.3 Experimental Results

**Overall Performance:** Fine-tuned mT5 models outperformed most language pairs and experimental settings. Notably, the GPT-3.5 and GPT-4 models demonstrated significant improvements in few-shot scenarios, particularly highlighting their effective adaptation in many-to-one settings, where they summarize from various source languages into English. However, in the one-to-many setting, GPT-3.5, GPT-4, and Mistral-7B-Instruct-v0.2 showed limited performance gains in the one-shot scenario, and their performance either deteriorated or remained unimproved in the two-shot setting.

Moreover, there was no significant performance difference among GPT-3.5, GPT-4, and Mistral-7B-Instruct-v0.2 in the one-to-many setting, indicating the challenges associated with summarizing from English to low-resource languages. The Mistral-7B-Instruct-v0.2 model consistently underperformed compared to the fine-tuned mT5 and the GPT-3.5 and GPT-4 models across most language pairs and few-shot settings, suggesting that it struggles to effectively adapt to the cross-lingual summarization task with limited examples.

**Few-Shot Learning Impact:** The performance of GPT-3.5 and GPT-4 models competitively improved as the number of shots increased, showcasing their few-shot learning capabilities in cross-lingual summarization. The Mistral-7B-Instruct-v0.2 model exhibited performance gains up to the one-shot setting, with generally increasing ROUGE scores. However, in the two-shot setting, the model’s performance showed a decreasing trend, indicating that the benefits of few-shot learning may not consistently extend to higher numbers of shots for this open-source model. This highlights the challenges in applying open-source models to few-shot cross-lingual summarization tasks and suggests that further research is needed to optimize their performance in these settings.

**Analysis by Language Pair:** The many-to-one approach generally resulted in higher ROUGE scores than the one-to-many approach. This suggests that summarizing in English is relatively more straightforward than summarizing in other languages. However, the performance gap between the two approaches was more pronounced for the Mistral-7B-Instruct-v0.2, indicating its limited ability to generate summaries in non-English target languages compared to the other models. Notably, all models achieved ROUGE scores of 0 for the English to Pashto language pair across all few-shot settings (see Table 2). This result indicates that few-shot learning did not improve the models’ performance for this specific language pair. The English to Pashto results, where all models failed to generate meaningful summaries even with few-shot learning, underscore the limitations of current approaches in handling extremely low-resource language pairs. This finding emphasizes the need for further research in developing more effective few-shot learning strategies and investigating the transfer learning capabilities of LLMs for cross-lingual summarization in such challenging scenarios.

## 5 Conclusion

This study empirically analyzed the few-shot performance of LLMs in cross-lingual summarization tasks, focusing on low-resource languages using a direct prompting approach. We observed that LLMs demonstrated competitive performance improvements through few-shot learning compared to zero-shot setups particularly in the many-to-one XLS. But, we also demonstrated that there was no significant gain to LLMs in the one-to-many XLS. These findings underscore the need for further research in developing more effective few-shot learning strategies and architectures tailored to low-resource languages.

## Limitation

Our study conducts experiments on a limited number of low-resource languages and uses only ROUGE metrics to validate the systems’ performance. Future research should explore advanced few-shot learning techniques, such as meta-learning or prompt-tuning, and investigate the impact of pre-training objectives and architectures designed specifically for cross-lingual tasks. This could lead to developing more effective open-source models for low-resource cross-lingual summarization.

Despite these limitations, this research demonstrates the potential of large language models’ few-shot learning capabilities in low-resource cross-lingual summarization tasks and provides experimental validation for the proposed research directions. Further work is necessary to extend these findings to additional low-resource languages and advance the few-shot learning capabilities of open-source models like Mistral-7B-Instruct-v0.2.

## Acknowledgement

This research was supported by Institute for Information & Communications Technology Planning & Evaluation (IITP) through the Korea government (MSIT) under Grant No. 2021-0-01341 (Artificial Intelligence Graduate School Program (Chung-Ang University)).

## References

Yu Bai, Yang Gao, and Heyan Huang. 2021. [Cross-lingual abstractive summarization with limited parallel resources](#). *Preprint*, arXiv:2105.13648.

- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. 2023. *A multitask, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity*. *Preprint*, arXiv:2302.04023.
- Abhik Bhattacharjee, Tahmid Hasan, Wasi Uddin Ahmad, Yuan-Fang Li, Yong-Bin Kang, and Rifat Shahriyar. 2023. *Crosssum: Beyond english-centric cross-lingual summarization for 1,500+ language pairs*. *Preprint*, arXiv:2112.08804.
- Florian Boudin, Stéphane Huet, and Juan-Manuel Torres-Moreno. 2011. *A graph-based approach to cross-language multi-document summarization*. *Polibits*, 43:113–118.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. *Language models are few-shot learners*. *Preprint*, arXiv:2005.14165.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. *Sparks of artificial general intelligence: Early experiments with gpt-4*. *Preprint*, arXiv:2303.12712.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th ophile Gervet, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2024. *Mixtral of experts*. *Preprint*, arXiv:2401.04088.
- La Plateforme. Mistral AI | frontier AI in your hands. <https://mistral.ai/>. Accessed: 2023-05-10.
- Faisal Ladhak, Esin Durmus, Claire Cardie, and Kathleen McKeown. 2020. *Wikilingua: A new benchmark dataset for cross-lingual abstractive summarization*. *Preprint*, arXiv:2010.03093.
- Anton Leuski, Chin-Yew Lin, Liang Zhou, Ulrich Germann, Franz Josef Och, and Eduard Hovy. 2003. *Cross-lingual c\*st\*rd: English access to hindi information*. *ACM Transactions on Asian Language Information Processing*, 2(3):245–269.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin 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*, pages 7871–7880, Online. Association for Computational Linguistics.
- Peiyao Li, Zhengkun Zhang, Jun Wang, Liang Li, Adam Jatowt, and Zhenglu Yang. 2023. *ACROSS: An alignment-based framework for low-resource many-to-one cross-lingual summarization*. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2458–2472, Toronto, Canada. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. *ROUGE: A package for automatic evaluation of summaries*. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- OpenAI. 2023. *GPT-4 technical report*. ArXiv. Abs/2303.08774.
- C Or san and OA Chiorean. 2008. *Evaluation of a cross-lingual romanian-english multi-document summariser*.
- Jacob Parnell, Inigo Jauregi Unanue, and Massimo Piccardi. 2024. *Sumtra: A differentiable pipeline for few-shot cross-lingual summarization*. *Preprint*, arXiv:2403.13240.
- Rajvardhan Patil and Venkat Gudivada. 2024. *A review of current trends, techniques, and challenges in large language models (llms)*. *Applied Sciences*, 14(5).
- Chengwei Qin, Aston Zhang, Zhuosheng Zhang, Jiaao Chen, Michihiro Yasunaga, and Diyi Yang. 2023. *Is ChatGPT a general-purpose natural language processing task solver?* In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1339–1384, Singapore. Association for Computational Linguistics.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timoth e Lacroix, Baptiste Rozi re, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. *Llama: Open and efficient foundation language models*. *Preprint*, arXiv:2302.13971.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruiti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura,

- Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Xiaojun Wan, Huiying Li, and Jianguo Xiao. 2010. [Cross-language document summarization based on machine translation quality prediction](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 917–926, Uppsala, Sweden. Association for Computational Linguistics.
- 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](#). *Preprint*, arXiv:2302.14229.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Jingfeng Yang, Hongye Jin, Ruixiang Tang, Xiaotian Han, Qizhang Feng, Haoming Jiang, Bing Yin, and Xia Hu. 2023. [Harnessing the power of llms in practice: A survey on chatgpt and beyond](#). *Preprint*, arXiv:2304.13712.
- Junnan Zhu, Qian Wang, Yining Wang, Yu Zhou, Jiajun Zhang, Shaonan Wang, and Chengqing Zong. 2019. [Ncls: Neural cross-lingual summarization](#). *Preprint*, arXiv:1909.00156.