

# Harnessing the Power of LLMs in Practice: A Survey on ChatGPT and Beyond

JINGFENG YANG\*, Amazon, USA

HONGYE JIN\*, Texas A&M University, USA

RUIXIANG TANG\*, Rice University, USA

XIAOTIAN HAN\*, Texas A&M University, USA

QIZHANG FENG\*, Texas A&M University, USA

HAOMING JIANG, Amazon, USA

SHAOCHEN ZHONG, Rice University, USA

BING YIN, Amazon, USA

XIA HU, Rice University, USA

This paper presents a comprehensive and practical guide for practitioners and end-users working with Large Language Models (LLMs) in their downstream natural language processing (NLP) tasks. We provide discussions and insights into the usage of LLMs from the perspectives of models, data, and downstream tasks. Firstly, we offer an introduction and brief summary of current language models. Then, we discuss the influence of pre-training data, training data, and test data. Most importantly, we provide a detailed discussion about the use and non-use cases of large language models for various natural language processing tasks, such as knowledge-intensive tasks, traditional natural language understanding tasks, generation tasks, emergent abilities, and considerations for specific tasks. We present various use cases and non-use cases to illustrate the practical applications and limitations of LLMs in real-world scenarios. We also try to understand the importance of data and the specific challenges associated with each NLP task. Furthermore, we explore the impact of spurious biases on LLMs and delve into other essential considerations, such as efficiency, cost, and latency, to ensure a comprehensive understanding of deploying LLMs in practice. This comprehensive guide aims to provide researchers and practitioners with valuable insights and best practices for working with LLMs, thereby enabling the successful implementation of these models in a wide range of NLP tasks. A curated list of practical guide resources of LLMs, regularly updated, can be found at <https://github.com/Mooler0410/LLMsPracticalGuide>.

CCS Concepts: • **Computing methodologies** → **Natural language processing**; **Natural language generation**; **Machine translation**.

Additional Key Words and Phrases: Large Language Models, Neural Language Processing, Practical Guide, ChatGPT

## 1 INTRODUCTION

In recent years, the rapid development of large Language models has been revolutionizing the field of natural language processing [13, 131, 134]. These powerful models have shown great potential in addressing a variety of NLP tasks and real-world cases, ranging from natural language understanding (NLU) to generation tasks, even paving the way to Artificial General Intelligence (AGI). However, utilizing these models effectively and efficiently requires a practical understanding of their capabilities and limitations, as well as the characteristics of the data and tasks.

---

\*These authors contributed equally.

---

Authors' addresses: Jingfeng Yang, [jingfengyangpku@gmail.com](mailto:jingfengyangpku@gmail.com), Amazon, USA; Hongye Jin, [jhy0410@tamu.edu](mailto:jhy0410@tamu.edu), Texas A&M University, USA; Ruixiang Tang, [rt39@rice.edu](mailto:rt39@rice.edu), Rice University, USA; Xiaotian Han, [han@tamu.edu](mailto:han@tamu.edu), Texas A&M University, USA; Qizhang Feng, [qf31@tamu.edu](mailto:qf31@tamu.edu), Texas A&M University, USA; Haoming Jiang, [jhaoming@amazon.com](mailto:jhaoming@amazon.com), Amazon, USA; Shaochen Zhong, [hz88@rice.edu](mailto:hz88@rice.edu), Rice University, USA; Bing Yin, [alexbyin@amazon.com](mailto:alexbyin@amazon.com), Amazon, USA; Xia Hu, [xia.hu@rice.edu](mailto:xia.hu@rice.edu), Rice University, USA.

To provide a guide for practitioners and end-users, this work focuses on the practical aspects of working with LLMs in the real world and downstream NLP tasks. This guide aims to provide practical advice on why or why not to choose LLMs for a given task, as well as guidance on how to select the most suitable LLM, taking into account factors such as model sizes, computational requirements, and the availability of domain-specific pre-trained models. This work offers a thorough understanding of LLMs from a practical perspective, therefore, empowers practitioners and end-users with the practical knowledge needed to successfully leverage the power of LLMs for their own tasks.

Our work is structured as follows. First, our work offers a brief introduction to LLMs by discussing the most important models, such as GPT-style and BERT-style architectures. Then, we delve into the critical factors that influence model performance from the data perspective, including pre-training data, training/tuning data, and test data. Last and most importantly, we dive deep into various concrete tasks, offering insights into the applicability of LLMs for knowledge-intensive tasks, traditional NLU tasks, and generation tasks, along with the emergent abilities that these models possess and challenging real-world scenarios. We provide detailed examples to highlight both the successful use cases and the limitations of LLMs in practice.

To analyze the abilities of large language models, we compare them with fine-tuned models. As of present, there is no universally recognized definition for LLMs and fine-tuned models. With consideration to practical utility, in our article, the definitions of them are proposed as: LLMs are huge language models pretrained on large amounts of datasets without tuning on data for specific tasks; fine-tuned models are typically smaller language models which are also pretrained and then further tuned on a smaller, task-specific dataset to optimize their performance on that task<sup>1</sup>.

This work summarizes the following main practical guides for using LLMs:

- **Natural language understanding.** Employ the exceptional generalization ability of LLMs when facing out-of-distribution data or with very few training data.
- **Generation.** Utilize LLMs' capabilities to create coherent, contextually relevant, and high-quality texts and code for various applications.
- **Knowledge-intensive tasks.** Leverage the extensive knowledge stored in LLMs for tasks requiring domain-specific expertise or general world knowledge.
- **Reasoning ability.** Understand and harness the reasoning capabilities of LLMs to improve decision-making and problem-solving in various contexts.
- **Real world scenarios.** Take the advantages of LLMs in real-world scenarios due to their ability to handle noisy input, tackle unformalized tasks, and follow human instructions after alignment.

Recently, numerous surveys have delved into the evolution of Large Language Models (LLMs). For instance, [131] offers a review of the recent advancements in LLMs, including their foundational principles, key contributions, and predominant techniques. [19] specifically focuses on evaluating the performance of LLMs, while [48] provides an in-depth overview of the current understanding of reasoning capabilities within these models. Additionally, [99] explores the challenges and opportunities related to discerning text generated by LLMs. Despite the wealth of existing literature, most surveys are tailored for an audience of computer science researchers. There is a noticeable gap in the provision of practical guidance for practitioners and end-users who may not have specialized knowledge of LLMs but wish to leverage their capabilities for downstream tasks. This survey aims to fill that void, offering a hands-on guide for a broader audience interested in harnessing the power of LLMs for various NLP applications.

<sup>1</sup>From a practical standpoint, we consider models with less than 20B parameters to be models which can be fine-tuned. While it's possible to fine-tune even larger models like PlaM (540B), in reality, it can be quite challenging, particularly for academic research labs and small teams. Fine-tuning a model with 3B parameters can still be a daunting task for many individuals or organizations.

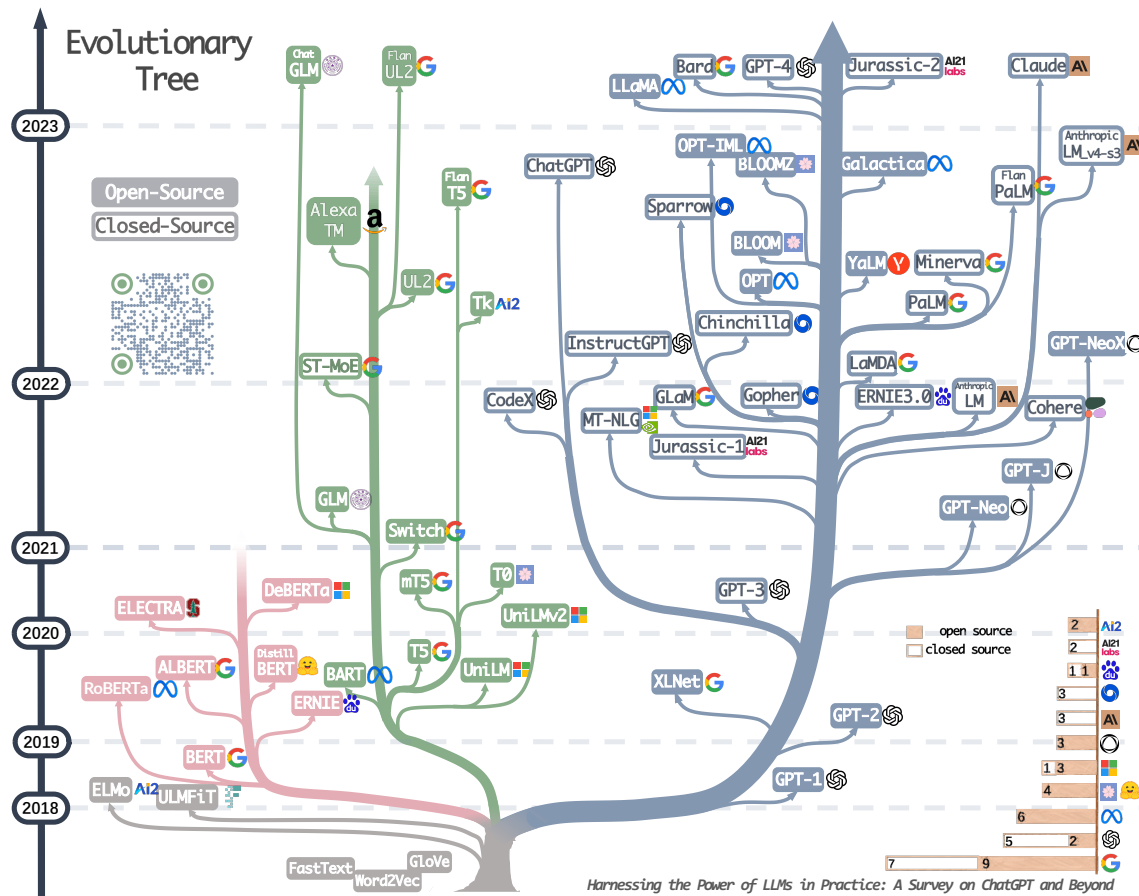


Fig. 1. The evolutionary tree of modern LLMs traces the development of language models in recent years and highlights some of the most well-known models. Models on the same branch have closer relationships. Transformer-based models are shown in non-grey colors: decoder-only models in the blue branch, encoder-only models in the pink branch, and encoder-decoder models in the green branch. The vertical position of the models on the timeline represents their release dates. Open-source models are represented by solid squares, while closed-source models are represented by hollow ones. The stacked bar plot in the bottom right corner shows the number of models from various companies and institutions. A dynamic version of this figure could be found by [this link](#).

## 2 PRACTICAL GUIDE FOR MODELS

This section provides a brief introduction to state-of-the-art LLMs. These models differ in their training strategies, model architectures, and use cases. To provide a clearer understanding of the LLM landscape, we categorize them into three types: non-causal attention language models, half-causal attention language models and causal attention language models<sup>2</sup>. In Figure 1, we show the detailed evolution process of language models. From the evolutionary tree, we make the following interesting observations:

<sup>2</sup>There are many ways to categorize LLMs. For example, LLMs can be categorized by the architecture, such as encoder-only, encoder-decoder or decoder-only. LLMs can also be categorized by the pretraining task, such as prefixLM and masked language modeling. It's hard to find a taxonomy which can include all LLMs. Some discussion may be helpful about the difference among LLMs [111].

- a) Decoder-only models have been gradually dominating the development of LLMs. At the early stage of LLMs development, **decoder-only** models were not as popular as **encoder-only** and **encoder-decoder** models. However, after 2021, with the introduction of game-changing LLMs - GPT-3, decoder-only models experienced a significant boom. Meanwhile, after the initial explosive growth brought about by BERT, encoder-only models gradually began to fade away.
- b) OpenAI consistently maintains its leadership position in LLM, both currently and potentially in the future. Other companies and institutions are struggling to catch up with OpenAI in developing models comparable to GPT-3 and the current GPT-4. This leadership position may be attributed to OpenAI's steadfast commitment to its technical path, even when it was not widely acknowledged initially.
- c) Meta contributes significantly to open-source LLMs and promotes research of LLMs. When considering contributions to the open-source community, particularly those related to LLMs, Meta stands out as one of the most generous commercial companies, as all the LLMs developed by Meta are open-sourced.
- d) LLMs exhibit a tendency towards closed-sourcing. In the early stages of LLM development (before 2020), the majority of models were open-sourced. However, with the introduction of GPT-3, companies have increasingly opted to close-source their models, such as PaLM, LaMDA, and GPT-4. Consequently, it has become more difficult for academic researchers to conduct experiments on LLM training. As a result, API-based research could become the predominant method in the academic community.
- e) Encoder-decoder models remain promising, as this type of architecture is still being actively explored, and most of them are open-sourced. Google has made substantial contributions to open-source encoder-decoder architectures. However, the flexibility and versatility of decoder-only models seem to make Google's insistence on this direction less promising.

## 2.1 BERT-style Language Models: Encoder-Decoder or Encoder-only

As natural language data is readily available and unsupervised training paradigms have been proposed to better utilize extremely large datasets, this motivates the unsupervised learning of natural language. Non-causal models are one of the earliest pretrained language models that sparked discussions. They are encoder-only models and adapt the approach of predicting masked words in a sentence while considering the surrounding context. This training paradigm is known as the masked language modeling. As the name suggests, for such models, all input tokens are visible to each other in the attention mechanism which does not follow the casualty of language. Notable examples of non-causal models include BERT [29] and RoBERTa [70]. Non-causal models have shown good performance in the field of natural language understanding.

## 2.2 Decode-only models: GPT-style language models

Although language models are typically task-agnostic in architecture, these methods require fine-tuning on datasets of the specific downstream task. Researchers found that scaling up language models significantly improves the few-shot, even zero-shot performance [17]. The most successful models for better few-shot and zero-show performance are causal attention language models. They are decoder-only and are trained by generating the next word in a sequence given the preceding words. For each token, only tokens before it is visible to this token in the attention mechanism following the casualty of human speaking. The GPT series belong to this category. They are also called autoregressive models. These models have been widely used for downstream tasks such as text generation and question answering. Examples

of autoregressive language Models include GPT-3 [17], OPT [129], PaLM [23], and BLOOM [95]. The game changer, GPT-3, for the first time, demonstrated reasonable few-/zero-shot performance via prompting and in-context learning.

### 2.3 Hybrid models

Beside the standard encode-decode model, encoder-only and decoder-only models, there are a series of works that employ hybrid pretraining objectives such as prefixLM [8, 89]. These models integrate aspects of both decoder and encoder models. For these models, the input text is divided into two parts when using the attention mechanism. The first part mirrors the encoder section in encoder-decoder models or the prefix segment in models using a prefixLM objective. However, the second part aligns with the method of decoder-only models, in which each token only accesses information from the tokens that came before it. The primary goal of these hybrid models is to boost the generative abilities of decoder-only models. Many popular models including T5 [89], UniLMv2 [8], and BART [62], belong to this category.

### 2.4 The Architecture of LLMs

The architecture of Language Models varies depending on their category—BERT-style (Encoder-Decoder or Encoder-only), GPT-style (Decode-only), or Hybrid models. In this section, we briefly introduce the architectures of the recent LLMs as follows:

- **Encoder-only or Encoder-Decoder models** introduced the concept of bidirectional pretraining. These models, such as BERT and RoBERTa, consist of an encoder-only architecture. In the case of encoder-only models, they use a masked language modeling training paradigm. This involves predicting masked words within a sentence while considering the surrounding context. Unlike decoder-only models, all input tokens in encoder-only models are visible to each other in the attention mechanism, disregarding the causality of language. On the other hand, encoder-decoder models use the encoder to process the input data and compress the information into a context-rich representation, and the decoder to generate the output sequence from this representation.
- **Decode-only models**, exemplified by the GPT series, are decoder-only architectures trained using autoregressive language modeling. Autoregressive training involves generating the next word in a sequence given the preceding words. Tokens in these models can only attend to tokens that precede them in the input sequence, adhering to the causal structure of language. GPT models have demonstrated impressive few-shot and zero-shot performance in tasks like text generation and question answering, especially when scaled up.

These three categories represent the main architectural paradigms of LLMs and have significantly contributed to the advancements in natural language understanding and generation. Researchers continue to explore and refine these architectures, pushing the boundaries of what LLMs can achieve in various language-related tasks.

When selecting an appropriate LLM architecture for the downstream tasks, several considerations come into play. If the task relies on capturing bidirectional contexts, such as Named Entity Recognition or document classification, then Encoder-only models like BERT and RoBERTa may be more suitable. On the other hand, if the task focus is on generating text like translation or summarization, then Decoder-only, or Encoder-Decoder models like GPT-series and BART-series could be more beneficial. Additionally, the model size is another crucial factor; larger models generally perform better but come at the cost of computational resources and latency. Therefore, balancing performance and resource availability is essential. Lastly, always consider the nature of the specific task, as specialized architectures might offer benefits over general-purpose models.

### 3 PRACTICAL GUIDE FOR DATA

In this section, we'll be discussing the critical role that data plays in selecting appropriate models for downstream tasks. The impact of data on the models' effectiveness starts during the pre-training stage and continues through to the training and inference stages.

#### Remark 1

- (1) LLMs generalize better than fine-tuned models in downstream tasks facing out-of-distribution data, such as adversarial examples and domain shifts. Evidenced by the comparison-based empirical results in related literature like [36, 88, 110], as well as the generalization improvement gained through alignment techniques like Reinforcement Learning from Human Feedback (RLHF) [83]. More in Section 3.3.
- (2) LLMs are preferable to fine-tuned models when working with limited annotated data, and both can be reasonable choices when abundant annotated data is available, depending on specific task requirements. This is evidenced by LLMs impressive zero-shot and few-shot learning ability [17, 124], as well as the fact that in-context learning does not require weight update, which mechanically prevents unfavorable fine-tuning-related behavior like catastrophic forgetting. More in Section 3.2.
- (3) It's advisable to choose models pre-trained on fields of data that are similar to downstream tasks. We investigate this data alignment on various popular LLMs in Section 3.1 and find preferable results when such alignment exists.

#### 3.1 Pre-training data

Pre-training data plays a pivotal role in the development of large language models. As the foundation of remarkable capabilities [4, 52] of LLMs, the quality, quantitative, and diversity of pre-training data influence the performance of LLMs significantly [127]. The commonly used pre-training data consists of a myriad of text sources, including books, articles, websites, and codes. The data is carefully curated to ensure a comprehensive representation of human knowledge, linguistic nuances, and cultural perspectives. The importance of pretraining data lies in its capacity to inform the language model with a rich understanding of word knowledge, grammar, syntax, and semantics, as well as the ability to recognize context and generate coherent responses. Recent studies indicate that pre-training data quality is a crucial factor affecting model performance. For example, a study by Gunasekar et al. [41] demonstrated that a Transformer-based model with only 1.3 billion parameters could achieve commendable performance on code generation tasks when trained on 6 billion tokens of 'textbook quality' web data and an additional 1 billion tokens generated synthetically using GPT-3.5. The diversity of pretraining data also plays a crucial role in shaping the model's performance, and the selection of LLMs highly depends on the components of the pretraining data. For example, PaLM [23] and BLOOM [95] excel in multilingual tasks and machine translation with an abundance of multilingual pretraining data. Moreover, PaLM's performance in Question Answering tasks is enhanced by incorporating a considerable amount of social media conversations and Books corpus [23]. Likewise, the code execution and code completion capabilities of GPT-3.5 (code-davinci-002) are amplified by the integration of code data in its pretraining dataset. In summary, when end-users aim to directly employ pre-trained models for their specific downstream tasks, it is recommended to select a model that has been pre-trained on data from a similar domain. If end-users intend to assemble their own pre-training data and train a custom model, both the quality and quantity of the data become critical factors that substantially influence the model's performance.

### 3.2 Finetuning data

When deploying a model for downstream tasks, it is essential to consider three primary scenarios based on the availability of annotated data: zero, few, and abundant. In this section, we provide a succinct overview of the appropriate models to employ for each scenario.

**Zero annotated data:** In scenarios where annotated data is unavailable, utilizing LLMs in a zero-shot setting proves to be the most suitable approach. LLMs have been shown to outperform previous zero-shot methods [124]. Additionally, the absence of a parameter update process ensures that catastrophic forgetting [54] is avoided since the language model parameters remain unaltered.

**Few annotated data:** In this case, the few-shot examples are directly incorporated in the input prompt of LLMs, which is also called in-context learning, and these examples can effectively guide LLMs to generalize to the task. As reported in [17], one-shot and few-shot performance make significant gains, even matching the performance of the SOTA fine-tuned open-domain models. And LLMs' zero/few-shot ability can be improved further by scaling [17]. Alternatively, some few-shot learning methods are invented to enhance fine-tuned models, such as meta-learning [59] or transfer learning [92]. However, performance might be inferior compared to using LLMs due to fine-tuned models' smaller scale and overfitting.

**Abundant annotated data:** With a substantial amount of annotated data for a particular task available, both fine-tuned models and LLMs can be considered. In most cases, fine-tuning the model can fit the data pretty well. Although, LLMs can be used to meet some constraints such as privacy [101]. In this scenario, the choice between using a fine-tuned model or a LLM is task-specific and also depends on many factors, including desired performance, computational resources, and deployment constraints.

In a brief summary: LLMs are more versatile w.r.t. data availability, while fine-tuned models can be considered with abundant annotated data.

### 3.3 Test data/user data

When deploying LLMs for downstream tasks, we often face challenges stemming from distributional differences between the test/user data and that of the training data. These disparities may encompass domain shifts [135], out-of-distribution variations [33], or even adversarial examples [88]. Such challenges significantly hinder fine-tuned models' effectiveness in real-world applications. They fit into a specific distribution and have a poor ability to generalize to OOD data. However, LLMs perform quite well facing such scenarios because they do not have an explicit fitting process. Moreover, recent advancements have further enhanced the ability of language models in this regard. The Reinforcement Learning from Human Feedback (RLHF) method has notably enhanced LLMs' generalization capabilities [83]. For example, InstructGPT demonstrates proficiency in following various instructions for a wide range of tasks and occasionally complying with instructions in different languages, even though such instructions are scarce. Similarly, ChatGPT exhibits consistent advantages on most adversarial and out-of-distribution (OOD) classification and translation tasks [110]. Its superiority in understanding dialogue-related texts led to an impressive performance on the DDXPlus dataset [104], a medical diagnosis dataset designed for OOD evaluation. In summary, LLMs generalize better than fine-tuned models in the most downstream tasks.

---

<sup>2</sup>As we mention in Section 1, LLMs are pretrained on large and diverse datasets without fine-tuning, while fine-tuned models are typically pretrained on a large dataset and then further trained on a smaller, task-specific dataset to optimize their performance on that task.

## 4 PRACTICAL GUIDE FOR NLP TASKS

In recent times, LLMs are predominantly generative models. Executing a typical NLP task with LLMs differs from the approach with traditional models. The latter often transforms labels into somewhat arbitrary integers. For LLMs, both the labels and input texts are first represented using natural language descriptions. The LLMs then process these natural language inputs and generate predictions, such as the label name, directly.

When utilizing LLMs, certain practical strategies can enhance their capabilities. Two of the most commonly employed are Chain-of-Thought (CoT) and In-context Learning (ICL):

**In-context Learning.** The in-context learning (ICL) paradigm is widely adopted with LLMs [17, 32, 73]. In ICL, LLMs make predictions based on contexts enhanced with a few sample examples. An LLM in ICL mode receives a set of examples and a query related to a specific task. This data is then presented as examples in a prompt to the model. Unlike traditional supervised learning, which updates model parameters during training, ICL does not alter parameters. It relies on the pretrained capabilities of LLMs, trusting them to identify patterns in the demonstrations and predict accurately. ICL can markedly improve LLM performance across various tasks. However, while ICL shows promise, several aspects remain elusive. Its efficacy can vary based on factors like the prompting template or the choice and order of in-context examples. Furthermore, the inner workings of ICL are not fully understood, with only a handful of studies investigating its mechanics [43, 96].

**Chain-of-Thought.** Despite LLMs' prowess in handling numerous tasks, they sometimes require augmentation for intricate reasoning tasks [24, 117]. The Chain-of-Thought (CoT) prompting method was devised to bolster LLM reasoning abilities. CoT demands that LLMs emulate a sequential, logical progression reminiscent of human thought processes. Using a few-shot approach, CoT supplies the LLM with task-solving steps in prompts, guiding it through a systematic reasoning process. Rather than offering a direct answer, the model is urged to methodically work through the problem, delineating logical steps, similar to human problem-solving techniques. In the zero-shot scenario, CoT merely instructs the model to approach the problem by "thinking through it step by step." CoT can enhance a model's proficiency in intricate reasoning tasks, yielding more precise and refined responses. It also provides a deeper insight into the model's cognitive processes. This concept has been expanded upon, leading to the "X-of-Thought" (XoT) methodology [10, 123], which encompasses a vast array of methods related to sequential reasoning in AI.

Following this understanding of how LLMs function in NLP tasks, we will next we discuss the specific scenarios where they are particularly effective (the use case), and contrastingly, situations where their application might not be the most optimal choice (the no use case).

### 4.1 Traditional NLU tasks

Traditional NLU tasks are some fundamental tasks in NLP including text classification, named entity recognition (NER), entailment prediction, and so on. Many of them are designed to serve as intermediate steps in larger AI systems, such as NER for knowledge graph construction.

#### Remark 2

Fine-tuned models generally are a better choice than LLMs in traditional NLU tasks, but LLMs can provide help while requiring strong generalization ability. We demonstrate how fine-tuned models are delivering superior performance to LLMs across such tasks in Section 4.1.1 and point out a few generalization-required NLU tasks where LLMs are preferable in Section 4.1.2.

*4.1.1 No use case.* In most natural language understanding tasks, such as tasks in GLUE[108] and SuperGLUE[107], fine-tuned models still have better performance, if such tasks come with rich well-annotated data and contain very few out-of-distribution examples on test sets. For different tasks and datasets, the gap between small fine-tuned models and LLMs varies.

In text classification, on most datasets, LLMs perform slightly worse than fine-tuned models. For sentiment analysis, such as on IMDB [75] and SST [97], fine-tuned models and LLMs perform equally well. For toxicity detection, which is another iconic text classification task, the gap is much larger. All LLMs cannot perform well on this task, and on CivilComments [14] even the best one is only better than random guessing [64]. On the other hand, most popular fine-tuned models can obtain much better performance [34]. and the Perspective API <sup>3</sup> is still one of the best for detecting toxicity. This API is powered by a multilingual BERT-based model, which is tuned on publicly available toxicity data and several smaller single-language CNNs distilled from this model. This might be due to the fact that toxicity is defined by subtle nuances in linguistic expressions, and large language models are unable to accurately comprehend this task solely based on the provided input.

The trend of performance gaps is similar in some other tasks. For natural language inference (NLI) tasks, on most datasets, such as on RTE [108] and SNLI [15], fine-tuned models perform better than LLMs, while on some data such as CB [107], LLMs have obtained comparable performance with fine-tuned models [23]. For question answering (QA), on SQuADv2 [90], QuAC [22] and many other datasets, fine-tuned models have superior performance, while on CoQA [91], LLMs perform as well as fine-tuned models [23].

In information retrieval (IR) tasks, LLMs are not widely exploited yet. One major reason is that IR tasks are fundamentally different from others. There’s no natural way to transform the thousands of candidate texts into a few/zero-shot form which is required by LLMs. The existing evaluation results on MS MARCO(regular/TREC) [79] show that methods based on fine-tuned models have better performance [64]. In this evaluation, the LLMs rank passages in an unorthodox way, which requires the LLMs to produce probabilities for passages one by one.

For some low-level intermediate tasks, which are not intended for regular users but rather for high-level tasks, such as named entity recognition (NER) and dependency parsing, there’s not enough result coming from LLMs, because the most current evaluation of LLMs focuses on practical tasks. According to available evaluation results, for the NER task, CoNLL03 [93] is still a challenge for LLMs [87], where the performance of fine-tuned models is around as twice as LLMs. These intermediate tasks may vanish soon because LLMs can take over high-level tasks without the help of those intermediate tasks (e.g. dependency parsing for coding tasks; NER for some text generation tasks).

In brief, for most traditional NLU tasks, a fine-tuned model is a better choice in terms of the performance on benchmark datasets and the computational cost. The scale of LLMs is usually 10× or even 100× larger than fine-tuned models. One possible cause for the inferior performance of LLMs on certain tasks can be the design of instructions/prompts. Transforming input from tasks like IR and sentence labeling into a few/zero-short instruction form is non-trivial. There may be better ways to adapt language models to traditional NLP tasks in the future. On the other hand, the upper limit of capabilities of fine-tuned models is not reached, and some methods like FLAN-tuning [72] can further boost the performance on NLU tasks. Another interesting finding is that on NLU tasks, after fine-tuning, masked language models, like T5[89], are better than most auto-regressive language models at the same scale, while some recent results imply that this gap can be bridged by scaling[23].

*4.1.2 Use case.* However, there are still some NLU tasks suitable for LLMs.

<sup>3</sup><https://perspectiveapi.com>

One of the representative tasks is miscellaneous text classification [64]. In contrast to classic domain-specific text classification tasks such as sentiment analysis, miscellaneous text classification deals with a diverse range of topics and categories that may not have a clear or strong relationship with one another. It's closer to real-world cases and hard to be formatted for using fine-tuned models. Another is the Adversarial NLI (ANLI)[80]. It is a difficult dataset composed of adversarially mined natural language inference questions in three rounds (R1, R2, and R3). LLMs have shown superior performance on ANLI, especially on the R3 and R2. Both examples demonstrate the exceptional ability of LLMs to generalize well on out-of-distribution and sparsely annotated data in traditional NLP tasks, surpassing that of fine-tuned models. We've discussed this in the section above 3.3.

## 4.2 Generation tasks

Generation tasks broadly encompasses two major categories of tasks, with the goal of creating coherent, meaningful, and contextually appropriate sequences of symbols. The first type focuses on converting input texts into new symbol sequences, as exemplified by tasks like paragraph summarization and machine translation. The second type, "open-ended" generation, aims to generate text or symbols from scratch to accurately match input descriptions such as crafting emails, composing news articles, creating fictional stories, and writing code.

### Remark 3

Due to their strong generation ability and creativity, LLMs show superiority at most generation tasks. We demonstrate the superiority of LLMs on various popular generative benchmarks like summarization, translation, coding, and even just some open-end generative tasks in Section 4.2.1. We also investigate some generative benchmarks where the LLMs fall short in Section 4.2.2, which generally happen due to extreme resource/data limitation.

*4.2.1 Use case.* Generation tasks require models to have a comprehensive understanding of the input contents or requirements and a certain level of creativity. This is what LLMs excel at.

For summarization tasks, although LLMs do not have an obvious advantage over fine-tuned models under traditional automatic evaluation metrics, such as ROUGE [65], human evaluation results indicate that humans tend to prefer the results generated by LLMs [40, 130] compared to that of fine-tuned models. For example, on CNN/DailyMail [77] and XSUM [78], fine-tuned models like Brio [71] and Pegasus [128] have much better performance than any LLMs w.r.t. ROUGE, but LLMs like OPT [129] perform far better in human evaluation considering all aspects including faithfulness, coherence, and relevance [130]. This demonstrates the superiority of LLMs in summarization tasks. On the other hand, it implies that current summarization benchmarks don't contain summaries with high quality or the automatic metrics are not proper for the evaluation of summarization.

In machine translation (MT), LLMs can perform competent translation, although the average performance is slightly worse than some commercial translation tools [50] considering some automatic metrics like BLEU[85]. LLMs are particularly good at translating some low-resource language texts to English texts, such as in the Romanian-English translation of WMT'16 [12], zero-shot or few-shot LLMs can perform better than SOTA fine-tuned model[23]. This is mainly due to the fact that English resources compose the main part of the pre-training data. BLOOM [95] is pre-trained on more multi-lingual data, leading to better translation quality in both rich-resource and low-resource translation. Another interesting finding is that BLOOM achieves good translation quality among Romance languages, even for translation from Galician, which is not included in the pre-training data. One reasonable explanation is that texts from

some languages in the same language group can help the LLMs learn more from the similarity. If more multi-lingual texts can be added to the pre-training data, the translation capability may be improved further.

Additionally, LLMs are highly skilled in open-ended generations. One example is that the news articles generated by LLMs are almost indistinguishable from real news articles by humans [17]. LLMs are remarkably adept at code synthesis as well. Either for text-code generation, such as HumanEval [20] and MBPP [6], or for code repairing, such as DeepFix [42], LLMs can perform pretty well. GPT-4 can even pass 25% problems in Leetcode, which are not trivial for most human coders [82]. With training on more code data, the coding capability of LLMs can be improved further [23]. While performing well on such tasks, the codes generated by LLMs should be tested carefully to figure out any subtle bugs, which is one of the main challenges for applying LLMs in code synthesis.

**4.2.2 No use case.** Fine-tuned models, such as DeltaLM+Zcode [120], still perform best on most rich-resource translation and extremely low-resource translation tasks. In rich resource machine translation, fine-tuned models slightly outperform LLMs [23, 95]. And in extremely low-resource machine translation, such as English-Kazakh translation, fine-tuned models significantly perform better than LLMs.

### 4.3 Knowledge-intensive tasks

Knowledge-intensive NLP tasks refer to a category of tasks that have a strong reliance on background knowledge, domain-specific expertise, or general real-world knowledge. These tasks go beyond simple pattern recognition or syntax analysis. And they are highly dependent on memorization and proper utilization of knowledge about specific entities, events, and common sense of our real world.

#### Remark 4

- (1) LLMs excel at knowledge-intensive tasks due to their massive real-world knowledge. Evidenced by their dominating performance on various general Question-Answering (QA) benchmarks as showcased in Section 4.3.1.
- (2) LLMs struggle when the knowledge requirements do not match their learned knowledge, or when they face tasks that only require contextual knowledge, in which case fine-tuned models can work as well as LLMs. We walk through some specific examples in this category in Section 4.3.2 and point out that this issue can be reasonably mitigated with the help of knowledge retrieval pipelines.

**4.3.1 Use case.** With billions of tokens and parameters, LLMs have more real-world knowledge than fine-tuned models.

Closed-book question-answering tasks require the model to answer a given question about factual knowledge without any external information. It does require the memorization of real-world knowledge in the model. LLMs perform better on nearly all datasets, such as on NaturalQuestions [57], WebQuestions [9], and TriviaQA [51]. On TriviaQA, even zero-shot LLMs is still much better [23].

The massive multitask language understanding (MMLU) [44] is also highly knowledge-intensive. It contains multiple-choice questions spanning over 57 different subjects and requires general knowledge of the model. It's pretty challenging even for LLMs, although the newly released GPT-4 [82] outperforms existing models by a considerable margin in English with a satisfactory 86.5% accuracy.

Also, some tasks in Big-bench[98], which are designed to probe LLMs and extrapolate their future capabilities, heavily relied on the memorization of real-world knowledge. In such tasks, the performance of some LLMs is better than the average level of humans, and even comparable to the best human performance. For example, the task *Hindu\_knowledge* requires models to give facts about Hindu mythology, *Periodic Elements* require the capability of predicting the element

name from the periodic table and *Physics* tests the physics knowledge of models by asking for the formula needed to solve a given physics problem.

4.3.2 *No use case.* There are some other tasks requiring knowledge different from that learned by LLMs. The required knowledge is not that learned by LLMs about the real world. In such tasks, LLMs are not notably superior.

Some tasks only require the model to capture the self-contained knowledge in the contexts. The knowledge in the contexts from the input is enough for the model to make predictions. For these tasks, small fine-tuned models can work pretty well. One such task is machine reading comprehension (MRC). An MRC task provides several paragraphs and requires the model to predict the answer to questions based on these paragraphs. We’ve discussed MRC in the previous section because it’s also a traditional NLU task.

Another scenario is that the knowledge within LLMs about real world is useless to the task, or even the required knowledge is counterfactual to the real world. As a result, the LLMs cannot work well on such tasks. In some cases, inconsistent knowledge may even make the LLMs worse than random guessing. For example, in Big-Bench, the Mnist ascii task requires the model to tell the digit represented by an ASCII art. The capability required by this task is nothing about real-world knowledge. Also, in the Inverse Scaling Phenomenon competition [76], the task *redefine math* redefines a common symbol and requires the model to choose between the original meaning and the meaning derived from the redefinition. What it requires contrasts with the LLMs’ knowledge, LLMs even perform worse than random guessing.

As an alternative to real-world knowledge in LLMs, access to extra knowledge is allowed, and models can thus get enough knowledge for a task via retrieval augmentation. The basic idea of retrieval augmentation is to add an extra information retrieval step prior to making predictions, in which, some useful texts related to the task will be retrieved from a large corpus. Then, the model will make predictions based on both the input contexts and the retrieved texts. With retrieved additional information, the closed-book task can become "open-book". In such a scenario, fine-tuned models are pretty good with much smaller sizes, because the required knowledge can be obtained by retrieving. For example, on NaturalQuestions [57], with extra corpus, retrieval augmented models [49, 53] are better than other methods.

#### 4.4 Abilities Regarding Scaling

Scaling of LLMs (e.g. parameters, training computation, etc.) can greatly empower pretrained language models. With the model scaling up, a model generally becomes more capable in a range of tasks. Reflected in some metrics, the performance shows a power-law relationship with the model scale. For example, the cross-entropy loss which is used to measure the performance for language modeling decreases linearly with the exponential increase in the model scale, which is also called 'scaling-law' [45, 52]. For some crucial abilities, such as reasoning, scaling the model has gradually transformed these abilities from a very low state to a usable state, and even approaching human capabilities. In this section, we provide an overview of the usage of LLMs in terms of the abilities and behaviors of LLMs along with scaling.

**Remark 5**

- (1) With the exponential increase of model scales, LLMs become especially capable of reasoning like arithmetic reasoning and commonsense reasoning. We give a walk-through of some popular reasoning benchmarks with respect to both reasoning categories in Section 4.4.1, as well as some techniques that may help on this school of tasks, like Chain-of-Thought (CoT) prompting [117].
- (2) Emergent abilities become serendipity for uses that arise as LLMs scale up, such as ability in word manipulation and logical ability. We discuss the definition of emergent abilities and specific task examples in Section 4.4.2.
- (3) In many cases, performance does not steadily improve with scaling due to the limited understanding of how large language models' abilities change as they scale up. We discuss such *inverse-scaling* and *U-shaped phenomena* in Section 4.4.3, where larger models are not always preferable.

**4.4.1 Use Case with Reasoning.** Reasoning, which involves making sense of information, drawing inferences, and making decisions, is one of the essential aspects of human intelligence. It is challenging for NLP. Many existing reasoning tasks can be classified into commonsense reasoning and arithmetic reasoning.

**Arithmetic reasoning/problem solving.** The arithmetic reasoning capability of LLMs benefits greatly from the scaling of model size [48]. For GPT-3, the ability of two-digit addition only becomes apparent when the number of parameters exceeds 13B [17]. Tasks to test arithmetic reasoning are trivial for humans and designed to challenge the capability of transferring natural language into mathematical symbols and multi-step inference. On GSM8k [27], SVAMP [86] and AQuA [66], LLMs, as generalists, have competitive performance with most methods which have task-specific designs. And GPT-4 overperforms any other methods [82], even some huge models particularly tuned for arithmetic problems [106]. Nevertheless, it should be noted that, without the intervention of external tools, LLMs may occasionally make mistakes in performing basic calculations, although chain-of-thought (CoT) prompting [117] can significantly improve LLMs' ability in calculations.

**Commonsense reasoning.** Commonsense reasoning not only requires LLMs to remember factual knowledge but also requires LLMs to do several inference steps about the facts. Commonsense reasoning increases gradually with the growth of model size. Compared to fine-tuned models, LLMs keep the superiority on most datasets, such as StrategyQA [38] and ARC-C [26]. Especially on ARC-C, which contains difficult questions in science exams from grade 3 to grade 9, GPT-4 has been close to the performance of 100% (96.3%) [82].

**4.4.2 Use Cases with Emergent Abilities.** Scaling of models also endows the model with some unprecedented, fantastic abilities that go beyond the power-law rule. These abilities are called "emergent ability". As defined in [115], *emergent abilities of LLMs are abilities that are not present in smaller-scale models but are present in large-scale models*. This means such abilities cannot be predicted by extrapolating the performance improvements on smaller-scale models and the model suddenly gains good performance on some tasks once the scale exceeds a certain range. The emergent ability is typically unpredictable and surprising, leading to tasks that emerge randomly or unexpectedly. We examine concrete examples of the emergent abilities of LLMs and provide them as an important reference for deciding whether to leverage LLMs' emergent abilities.

Handling word manipulation is a typical emergent ability. It refers to the ability to learn symbolic manipulations, such as the reversed words [17], in which the model is given a word spelled backwards, and must output the original word. For example, GPT-3 [17] shows the emergent ability for word sorting, and word unscrambling tasks. PaLM [23] exhibits

the emergent ability on ASCII word recognition<sup>4</sup> and hyperbaton<sup>5</sup> task. The logical abilities of language models tend to emerge as the model scales up, such as logical deduction, logical sequence, and logic grid puzzles. Additionally, other tasks, such as advanced coding (e.g., auto debugging, code line description), and concept understanding (e.g., novel concepts, simple Turing concepts), are also use cases with the emergent abilities of large language models.

**4.4.3 No-Use Cases and Understanding.** Although in most cases, as discussed above, larger models bring better performance, there are still many exceptions that should be considered when choosing the appropriate model.

On certain tasks, with the size of LLMs increasing, the performance begins to decrease, such as Redefine-math: tests whether language models are able to work with common symbols when they are redefined to mean something else; Into-the-unknown: requires the model to choose which piece of information would help answer a question; Memo-trap: asks an LM to write a phrase in a way that starts like a famous quote but ends differently<sup>6</sup>. This is also called *Inverse Scaling Phenomenon*. Another interesting phenomenon observed in the scaling of LLMs is called the *U-shaped Phenomenon* [116]. As the name implies, This phenomenon refers to that as LLM size increases, their performance on certain tasks initially improves but then starts to decline before eventually improving again, such as on: Hindsight-neglect: it tests whether language models are able to assess whether a bet was worth taking based on its expected value; NegationQA: this task takes an existing multiple-choice dataset and negates a part of each question to see if language models are sensitive to negation; Quote-repetition: it asks models to repeat back sentences given in the prompt, with few-shot examples to help it recognize the task. Hence the risk of diminishing performance should be noted and if the task is similar to those we just discussed, careful consideration should be given to whether or not to use huge LLMs.

Gaining a deeper understanding of emergent abilities, inverse scaling phenomenon and U-shape phenomenon in LLMs is essential for advancing research in this field. In a certain sense, the U-shape phenomenon suggests that small-scale models and huge-scale models make predictions with different internal mechanisms. From this perspective, the U-shape phenomenon can be seen as a transformation of the inverse-scaling phenomenon due to some emergent abilities from sufficiently large models [116]. GPT-4 [82] exhibits a reversal of the inverse scaling phenomenon in some cases, such as on a task called Hindsight Neglect. The explanation for these behaviors of LLMs during scaling is still an open problem. Several hypotheses have been proposed. For emergent abilities, one explanation is that there may be multiple key steps for a task and the LLM cannot handle this task until it's large enough to handle every step, and another explanation is focused on the granularity of evaluation metrics [115]. For inverse-scaling phenomenon and u-shape phenomenon, the explanations mainly focus on the model's over-reliance on information from its prior rather than the input prompts, valid but misleading few-shot examples, and distracting easier tasks within a hard task [116].

## 4.5 Miscellaneous tasks

This section explores miscellaneous tasks which cannot be involved in previous discussions, to better understand LLMs' strengths and weaknesses.

<sup>4</sup>Asking models to identify the word displayed as ASCII art, [https://github.com/google/BIG-bench/tree/main/bigbench/benchmark\\_tasks/ascii\\_word\\_recognition](https://github.com/google/BIG-bench/tree/main/bigbench/benchmark_tasks/ascii_word_recognition)

<sup>5</sup>Asking models to choose the English sentence with adjectives in the "correct" order within two choices, [https://github.com/google/BIG-bench/tree/main/bigbench/benchmark\\_tasks/hyperbaton](https://github.com/google/BIG-bench/tree/main/bigbench/benchmark_tasks/hyperbaton)

<sup>6</sup>More such tasks include: modus-tollens, pattern-matching-suppression, prompt-injection, repetitive-algebra, and sig-figs. You can check them on: <https://github.com/inverse-scaling/prize>

**Remark 6**

- (1) Fine-tuned models or specified models still have their space in tasks that are far from LLMs' pretraining objectives and data, as indicated in Section 4.1.1 and Section 4.5.1, where the lack resource-alignment is one of the major cause.
- (2) LLMs are excellent at mimicking human, data annotation and generation. They can also be used for quality evaluation in NLP tasks and have bonuses like interpretability. We provide several successful examples of such LLM-based generation/evaluation or human-LLM collaborations in Section 4.5.2.

4.5.1 *No use case.* LLMs generally struggle with some tasks due to differences in objectives and training data.

Although LLMs have achieved remarkable success in various natural language processing tasks, their performance in regression tasks has been less impressive. For example, ChatGPT's performance on the GLUE STS-B dataset, which is a regression task evaluating sentence similarity, is inferior to a fine-tuned RoBERTa performance [133]. The Regression tasks typically involve predicting a continuous value rather than a discrete label, posing unique challenges for LLMs. One primary reason for their subpar performance is the inherent difference between the language modeling objective and the regression task objective. LLMs are designed to predict the next word in a sequence or generate coherent text, with their pre-training focused on capturing linguistic patterns and relationships. Consequently, their internal representations may not be well-suited for modeling continuous numerical outputs. Besides, LLMs have predominantly been trained on text data, focusing on capturing the intricacies of natural language processing. As a result, their performance on multimodal data, which involves handling multiple data types such as text, images, audio, video, actions, and robotics, remains largely unexplored. And fine-tuned multimodal models, like BErT[112] and PaLI [21], still dominate many tasks such as visual question answering (VQA) and image captioning. Nonetheless, the recently introduced GPT-4 [82] has taken the step in multimodal fusion, but there is still a lack of detailed evaluation of its capabilities.

4.5.2 *Use case.* LLMs are particularly suitable for certain tasks.

LLMs are good at mimicking humans, acting as a chatbot, and performing various kinds of tasks. The LLMs-powered ChatGPT<sup>7</sup> is surprising for its consistency, reliability, informativeness, and robustness during multiple utterances with humans. The human-feedback procedure plays an important role in acquiring such abilities

LLMs can both act as a good annotator and data generator for data augmentation, such as in[28, 30, 101, 125, 126]. Some LLMs have been found as good as human annotators [39] in some tasks. The collected texts from GPT-3.5 (text-davinci-003) have been used as human-like instruction-following demonstrations to train other language models [103].

LLMs can also be used for quality assessment on some NLG tasks, such as summarization and translation. On summarization tasks, GPT-4 as an evaluator achieves a higher correlation with humans than other methods with a large margin [69]. Some other evaluators based on LLMs [35, 55, 69, 109] also show good human alignment in more NLG tasks, especially compared with traditional metrics. But the LLM may have a bias towards the LLM-generated texts [69].

Also, as we discussed above, some abilities of LLMs bring bonuses in addition to performance improvement, such as interpretability. The CoT reasoning ability of LLMs can show how an LLM reaches the prediction, which is a good interpretation on the instance level, while it also improves performance.

<sup>7</sup><https://chat.openai.com>

#### 4.6 Real world "tasks"

In the last part of this section, we would like to discuss the usage of LLMs and fine-tuned models in real-world "tasks". We use the term "tasks" loosely, as real-world scenarios often lack well-formatted definitions like those found in academia. Many requests to models even cannot be treated as NLP tasks. Models face challenges in the real world from three perspectives:

- **Noisy/Unstructured input.** Real-world input comes from real-world non-experts. They have little knowledge about how to interact with the model or even cannot use texts fluently. As a result, real-world input data can be messy, containing typos, colloquialisms, and mixed languages, unlike those well-formed data used for pre-training or fine-tuning.
- **Tasks not formalized by academia.** In real-world scenarios, tasks are often ill-defined by academia and much more diverse than those in academic settings. Users frequently present queries or requests that do not fall neatly into predefined categories, and sometimes multiple tasks are in a single query.
- **Following users' instructions.** A user's request may contain multiple implicit intents (e.g. specific requirement to output format), or their desired predictions may be unclear without follow-up questions. Models need to understand user intents and provide outputs that align with those intents.

Essentially, these challenges in the real world come from that users' requests deviate significantly from the distribution of any NLP datasets designed for specific tasks. Public NLP datasets are not reflective of how the models are used [83].

##### Remark 7

LLMs are better suited to handle real-world scenarios compared to fine-tuned models. However, evaluating the effectiveness of models in the real world is still an open problem. In Section 4.6, we analyze why such tasks are challenging to solve, and oftentimes, even hard to properly define and evaluate. We also discuss why LLMs may have the competitive advantage on such tasks due to their impressive generalization power gained via pretraining on mass resources and tuning to provide (human) preferable outputs.

Handling such real-world scenarios requires coping with ambiguity, understanding context, and handling noisy input. Compared to fine-tuned models, LLMs are better equipped for this because they have been trained on diverse data sets that encompass various writing styles, languages, and domains. Essentially, LLMs has seen nearly the whole real-world text distribution during pretraining, making them well-suited for these scenarios. Fine-tuned models, on the other hand, are often tailored to specific, well-defined tasks and may struggle to adapt to new or unexpected user requests. They heavily rely on clear objectives and well-formed training data that specify the types of instructions the models should learn to follow. Fine-tuned models may struggle with noisy input due to their narrower focus on specific distributions and structured data. An additional system is often required as an assistant for fine-tuned models to process unstructured context, determine possible intents, and refine model responses accordingly. Note that, to enable LLMs to better respond to open-ended user inputs and instructions, diverse instruction fine-tuning is still typically required to better align with real-world user input distribution and human intents of output. RLHF is then used as an additional step after instruction finetuning to precisely adjust the model input and output distribution to achieve the same goal. Nevertheless, the major capability of LLMs still rooted in the pretraining stage.

Additionally, some mechanics such as instruction tuning [94, 114] and human alignment tuning [83] further boost the capabilities of LLMs to better comprehend and follow user instructions. These methods improve the model's ability to generate helpful, harmless, and honest responses while maintaining coherence and consistency [83, 94, 114]. While both methods can make LLMs better generalize to unseen tasks and instructions, it has been noticed that while human

labelers prefer models tuned for human alignment [83] to models tuned with instructions from public NLP tasks, such as FLAN [114] and T0 [94]. The reason may be similar to reasons for fine-tuned models' inferiority: public NLP tasks/datasets are designed for easy and automatic evaluation, and they can only cover a small part of real-world usage.

One of the main issues when it comes to real-world scenarios is how to evaluate whether the model is good or not. Without any formalized tasks or metrics, the evaluation of model effectiveness can only rely on feedback from human labelers. Considering the complexity and cost of human evaluation, there's no massive and systematic comparison between fine-tuned models and LLMs yet. Nevertheless, the huge success and popularity of LLMs such as chatGPT, have confirmed the superiority of LLMs to some extent.

#### 4.7 Exemplary products

In the sections above, we discussed the *use case* and *no use case* of various common tasks, aiming to provide practical guidance to an intended user of LLMs. However, building a sound LLM-powered product never stops at understanding the (a selective portion of) conceptual advantages and limitations of LLM under different task scenarios, as a successful product will often need to clearly define its function scope and target users, then optimize upon it. To facilitate such a practical initiative, here we provide a walkthrough of some impactful LLM-powered tools and products, serving as examples to inspire and guide future LLM-powered product developers.

The first example is the well-known ChatGPT developed by OpenAI<sup>8</sup>, power by its GPT model family as based models, ChatGPT attracts the most attention and started the current wave of LLM-powered products, for its amazing interactive capability, even on highly custom and out-of-distribution tasks [17]. ChatGPT was originally launched with GPT-3.5 as its backbone with a pure chatbot interface, but it then received various core model and wrapper functionality updates, e.g., GPT-3.5, GPT-3.5-turbo, and GPT-4 as optional backbone to optimize efficiency and inference cost [82]; and custom instructions<sup>9</sup> (storage of prewritten prompt instructions), code interpreter<sup>10</sup> (a sandbox Python environment), retrieval query<sup>11</sup>, and various other Plugins<sup>12</sup> (integrated APIs for expanded functionality), which aim to improve the usability of ChatGPT under repetitive tasks and expands its capability beyond a close-source text-only chatbot.

Building upon OpenAI inference APIs, various GPT-powered third-party services have gained significant traction for being able to provide custom solutions to problems that are not easily resolvable under the ChatGPT chatbot interface or other types of limitations (e.g., context window). One emblematic example is ChatPDF<sup>13</sup> and its open-sourced counterpart, pdfGPT<sup>14</sup>, which allows users to upload a reasonably long PDF and ask questions about it. ChatPDF is able to deliver impressive answers based on the input PDF document and provide reference to the original source. We'd like to note that ChatGPT is also able to take a PDF input and perform document-based Q&A upon the usage of certain plugins, though the user experience is not exactly ideal due to token limitation and the lack of interactive UI to jump the source document and answers. The wide custom adaption of tools like ChatPDF highlights an interesting notion: that a healthy amount of "assistant"-type LLM-powered products can be successfully built with a user experience-driven approach instead of competing in the arm-race of model capability that is prevalently found among academic scholars.

<sup>8</sup><https://chat.openai.com>

<sup>9</sup><https://openai.com/blog/custom-instructions-for-chatgpt>

<sup>10</sup><https://openai.com/blog/chatgpt-plugins#code-interpreter>

<sup>11</sup><https://github.com/openai/chatgpt-retrieval-plugin>

<sup>12</sup><https://openai.com/blog/chatgpt-plugins>

<sup>13</sup><https://www.chatpdf.com>

<sup>14</sup><https://github.com/bhaskatripathi/pdfGPT>

Like ChatGPT being a natural language assistant, GitHub Copilot<sup>15</sup> provides assistance under a coding context and can provide code completion and generation upon existing input or commented instructions. Yet, products like Google Cloud Student Success Service<sup>16</sup> aim to provide personalized tutoring experiences for individual students. On the creative side, we have Adobe Generative Fill<sup>17</sup>, a Vision-Language Model trained upon licensed in-house data, to provide a visual “fill” according to the surrounding content and user instructions. We have also seen products like MuseNet<sup>18</sup> to generative songs and lyrics following given instructions. The rise of such (partially) LLM-powered multimodal products showcased the popularity of general AIGC (artificial intelligence-generated content), with many opportunities and challenges involved (e.g., how to commercialize AI-generated art when requires pre-training).

Given the popularity of LLM-powered products, we are in no way able to provide an exhaustive walkthrough of such products, but we do hope the above mentions can be used as the breaking point and learning reference for aspired LLM-powered product enthusiasts. We would also like to quickly note the contribution of open-source projects, such as HuggingFaces [118], LLaMA [105], LangChain<sup>19</sup>, Promptify [84], AutoGPT<sup>20</sup>, and many others. These projects often serve as the foundation of open-sourced LLM-powered products or are exemplary attempts to expand the capability of LLM onto other fields and tasks. Such projects play many important roles in facilitating the democratic progress of the LLM community and pave the way for a more accessible and inclusive tech landscape.

#### 4.8 Summary

We summarize all discussions into a decision flow in Figure 2. When a user facing a task, this decision flow can be a guide for quick decision of using LLMs, such as GPT-4, or fine-tuning a smaller model for this task. For example, a task translating English to Chinese: → 1. it’s not related to mimicking human behaviors; → 2. it does not require scaling; → 3. it does not contain multiple tasks; → 4. it’s related to language modeling; → 5. English-Chinese data is rich; → 6. It’s not creative generation; → 7. it’s not a knowledge-intensive task → **fine-tuned models**.

## 5 OTHER CONSIDERATIONS

Despite LLMs are suitable for various downstream tasks, there are some other factors to consider, such as efficiency and trustworthiness. Our discussion of efficiency encompasses the training cost, inference latency, and parameter-efficient tuning strategies for LLMs. Meanwhile, our examination of trustworthiness includes robustness & calibration, fairness & biases, potential spurious correlations, and the safety challenges in LLMs.

<sup>15</sup><https://github.com/features/copilot>

<sup>16</sup><https://cloud.google.com/blog/topics/public-sector/new-google-cloud-student-success-services-help-educators-scale-individualized-learning>

<sup>17</sup><https://www.adobe.com/products/photoshop/ai.html>

<sup>18</sup><https://openai.com/research/musenet>

<sup>19</sup><https://github.com/langchain-ai/langchain>

<sup>20</sup><https://github.com/Significant-Gravitas/AutoGPT>

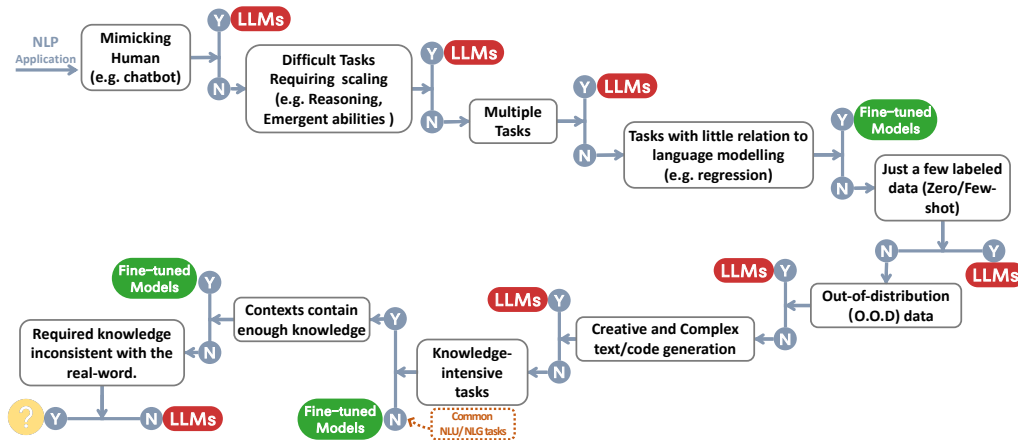


Fig. 2. The decision flow for choosing LLMs or fine-tuned models<sup>21</sup> for user’s NLP applications. The decision flow helps users assess whether their downstream NLP applications at hand meet specific conditions and, based on that evaluation, determine whether LLMs or fine-tuned models are the most suitable choice for their applications. During the decision process in the figure, **Y** means meeting the condition, and **N** means not meeting the condition. The yellow circle for **Y** of the last condition means there’s no model working well on this kind of application.

**Remark 8**

- (1) Light, local, fine-tuned models should be considered rather than LLMs, especially for those who are sensitive to the cost or have strict latency requirements. Parameter-Efficient tuning can be a viable option for model deployment and delivery.
- (2) The zero-shot approach of LLMs prohibits the learning of shortcuts from task-specific datasets, which is prevalent in fine-tuned models. Nevertheless, LLMs still demonstrate a degree of shortcut learning issues.
- (3) Safety concerns associated with LLMs should be given utmost importance as the potentially harmful or biased outputs, and hallucinations from LLMs can result in severe consequences. Some methods such as human feedback have shown promise in mitigating these problems.

**5.1 Efficiency**

In real-world deployment, performance, cost, and latency are all important considerations, not just the performance of the models. While some parameter-efficient methods have been developed, practitioners must balance efficiency with effectiveness in the practice.

*Cost.* LLMs have grown increasingly larger in recent years, with models such as GPT-1, GPT-2, and GPT-3 featuring 117 million, 1.5 billion, and 175 billion parameters, respectively. The cost of training an LLM is heavily influenced by its size, with estimates suggesting that training the 11B parameter variant of T5 costs well over \$1.3 million for a single run, while a single training run of GPT-3 175B requires \$4.6 million [2]. The energy consumption for training large models is equally impressive. The total energy consumption for training a transformer model with 6B parameters to completion is estimated to be around 103.5 MWh [31]. Google reports that training PaLM consumed about 3.4 GWh in about two months [5]. Furthermore, the dataset size also scales rapidly with the size of the model, with GPT-3 175B trained on 499 billion tokens [17]. Another key metric that reflects the computing cost is Flops, with GPT-3 175B requiring  $3.14 \times 10^{23}$  Flops, while a T5 11B model only requires  $3.30 \times 10^{22}$ , which is 10 times less. In addition to these costs, hardware requirements are also substantial. OpenAI has collaborated with Microsoft on a supercomputer hosted in the Microsoft Azure cloud, consisting of 285k CPU cores and 10k high-end GPUs to support the training of large models. For users of

the OpenAI API, pricing varies based on the model and usage, with options such as GPT-3.5-turbo charging \$0.002 per 1k tokens for chat service. However, for users who require custom models, training costs \$0.03 per 1k tokens, while usage costs \$0.12 per 1k tokens [3]. Therefore, for users who cannot afford such a large cost, such as small startups, individual users, etc., a small, fine-tuned model is a better and more reasonable choice.

*Latency.* Latency is a crucial factor to consider in real-world applications of LLMs. Inference time is a commonly used metric to measure latency, which is highly dependent on the model size, architecture, and token size. For instance, the inference time for the GPT-J 6B model is 0.077s, 0.203s, and 0.707s when the max token size is set to 2, 8, and 32, respectively [64]. Additionally, when the max token size is fixed at 32, the inference time for the InstructGPT model (davinci v2) is 1.969s. As LLMs are often too large to be run on a single user’s machine, companies provide LLM services via APIs. The API latency can vary depending on the user’s location, and the average latency of the OpenAI API service for a single request can range from a few hundred milliseconds to several seconds. In scenarios where high latency is not acceptable, large LLMs may not be appropriate. For example, scalability is critical in many information retrieval applications. To deploy information retrieval systems on the web, search engines require very efficient inference for systems to be useful. The idealized denoised inference time for the InstructGPT davinci v2 (175B\*) model is 0.21s per request (i.e., a query-passage pair to be scored), which is too slow for web search engines.

*Parameter-Efficient Tuning.* In practice, we may tune the model on some specific datasets. Parameter-Efficient Tuning (PET) is an efficient technique to tune a small portation of model parameters (or extra parameters) while freezing most parameters of the pre-trained LLMs. The main goal of PEFT is to greatly decrease the computational and storage costs while keeping the performance of the original models. The common techniques for PET are LoRA [46], Prefix Tuning [63], P-Tuning [67, 68]. As an illustration, the LoRA method maintains the weights of the pre-trained model and incorporates low-rank matrices into every layer of the Transformer architecture. This approach considerably minimizes the number of parameters that require training for subsequent tasks, thereby increasing overall efficiency. Alpaca-LoRA<sup>22</sup> proposes integrating Low-Rank Adaptation (LoRA) into LLaMA-Alpaca, which enables runs LLaMA within hours on a single RTX 4090. All these PFT methods can be helpful either for fine-tuning a model to a specific task or tuning LLMs to meet special requirements like human alignment.

## 5.2 Trustworthiness

Given that LLMs are now involved in sensitive areas such as healthcare, finance, and law, it is crucial to ensure that they are trustworthy and capable of producing reliable output.

*Robustness and Calibration.* The accuracy and robustness of the LLMs are shown to have a very strong correlation [64]. The models that have high accuracy on the scenario also have good robustness. However, the robustness of the zero-shot becomes worse after being tuned on extra application-specific tasks data [119]. This may due to overfitting, which leads to poor generalizability due to the extremely high complexity of the model and the limited training samples from downstream tasks [47]. In a similar vein, it has been observed that fine-tuning a model can result in significant miscalibrations, owing to over-parameterization [56]. Therefore, fine-tuned models may not be an optimal choice when robustness and calibration are critical considerations. However, human alignment has been found as a potential solution for enhancing model robustness. InstructGPT davinci v2 (175B\*) has been shown to outperform other models in terms

<sup>22</sup><https://github.com/tloen/alpaca-lora>

of robustness. On the other hand, achieving optimal calibration of the model depends on the scenario and adaptation procedure employed.

*Fairness and Bias.* LLMs have been shown to exhibit disparate treatment and impact, perpetuating societal biases and potentially leading to discrimination [11, 18, 60, 61]. To ensure fairness and equity for all users, it is crucial to address these issues in the development and deployment of NLP models. Disparities in performance between demographic groups can serve as an indicator of fairness problems. LLMs are particularly susceptible to fairness issues, as significant performance disparities have been observed across demographic categories such as dialect, religion, gender, and race [64]. However, research has shown that aligning models with human instructions can improve LLM performance regardless of their size, with the InstructGPT model (davinci v2) exhibiting smaller performance disparities than other LLMs [25].

*Spurious Biases.* The shortcut learning problem has been observed in various natural language understanding tasks under the pretraining and fine-tuning paradigm, where models heavily rely on spurious correlations between input and labels in the fine-tuning data for prediction [33, 37, 100, 102, 121, 122]. For example, in reading comprehension tasks, fine-tuned models tend to focus on the lexical matching of words between the question and the original passage, neglecting the intended reading comprehension task itself [58]. In contrast, large language models are not directly trained on fine-tuned datasets, which makes it less likely for them to learn shortcut features present in the fine-tuned dataset, thereby enhancing the model's generalization capabilities. However, LLMs are not infallible and may exhibit some shortcut learning during in-context learning. For example, recent preliminary studies have begun investigating the robustness of prompt-based methods in large-scale language models [102, 113, 132]. One such study evaluates the few-shot learning performance of GPT-3 on text classification and information extraction tasks [132], and reveal that the examined LLMs are susceptible to majority label bias and position bias, where they tend to predict answers based on the frequency or position of the answers in the training data. Moreover, these LLMs exhibit common token bias, favoring answers that are prevalent in their pre-training corpus. Recent studies show that this positional bias can be mitigated by selecting proper prompts [74]. In summary, while LLMs significantly reduce the shortcut learning problem prevalent in fine-tuned models, they still exhibit some shortcut learning issues and should be approached with caution when deploying them in downstream applications.

### 5.3 Safety challenges

LLMs have demonstrated their extremely strong capabilities in many areas such as reasoning, knowledge retention, and coding. As they become more powerful and human-like, their potential to influence people's opinions and actions in significant ways grows. As a result, some new safety challenges to our society should be considered and have caught lots of attention in recent works [81, 82].

*Hallucinations.* The potential for LLMs to "hallucinate," or generate nonsensical or untruthful content, can have significant negative impacts on the quality and reliability of information in various applications. As LLMs become increasingly convincing and believable, users may develop an overreliance on them and trust them to provide accurate information in areas with which they are somewhat familiar. This can be particularly dangerous if the model produces content that is entirely false or misleading, leading to incorrect decisions or actions taken based on that information. Such outcomes can have serious consequences in many domains, such as healthcare, finance, or public policy, where

the accuracy and reliability of information are critical. To mitigate these issues, reinforcement learning from human feedback (RLHF) is widely used [81, 83] and LLMs themselves have been integrated into the loop [81].

*Harmful content.* Due to the high coherence, quality, and plausibility of texts generated by LLMs, harmful contents from LLMs can cause significant harm, including hate speech, discrimination, incitement to violence, false narratives, and even social engineering attack. The implementation of safeguards to detect and correct those contents can be mitigation [99]. These LLMs can also have dual-use potential by providing required illicit information, leading to risks such as the proliferation of weapons [81] and even terrorism attack planning. It is crucial to ensure using these LLMs responsibly, with safeguards in place to prevent harm. Also, in existing work, feedback from humans plays an important role in getting rid of harmful outputs.

*Privacy.* LLMs can face serious security issues. An example is the issue of user privacy. It is reported that Samsung employees were using ChatGPT to process their work when they inadvertently leaked top-secret data, including the source code proper of the new program, internal meeting minutes related to the hardware, etc. The Italian data protection agency declared that OpenAI, the developer of ChatGPT, illicitly gathered personal user data, leading Italy to become the first government to prohibit ChatGPT over privacy concerns [1].

## 6 CHALLENGES IN PRACTICE

Despite the significance of LLMs, there are critical challenges of using LLMs in practice. In the following, we figure out some of them.

- **Evaluation of Proposed Models on Real-World “Datasets”.** The evaluation of deep learning models, for the most part, has been based on standard academic datasets, like ImageNet. Such datasets, despite their significance in the evolution of AI, have limitations in their ability to reflect real-world performance accurately. To truly gauge the capabilities and practical applicability of models, it is imperative to assess them using data that is diverse, complex, and mirrors real-world scenarios. As the scale of pre-training datasets expands, individual researchers face challenges in comprehensively reviewing and ensuring the quality of entire documents. Approximate duplicate data could adversely affect model performance. Effective filtering mechanisms, such as the use of MinHash algorithms, are becoming increasingly necessary to reduce redundancy. For pre-trained models fine-tuned on multiple tasks, an appropriate task-mix ratio is crucial. While appending task descriptions to each input-output pair is a common strategy, striking a balance between datasets remains a challenge. Moreover, while some open-source models try to emulate proprietary ones, significant capability gaps persist. Including data in training sets that are similar or related to evaluation test sets could inflate performance metrics, with models possibly memorizing test data. It’s worth noting that undetected personal identifiers, such as phone numbers or email addresses in pre-training datasets, raise severe privacy concerns. Future work should prioritize robust methods for data cleansing, minimizing biases, and ensuring privacy.
- **Challenges in Predicting Performance with Scaling.** The landscape of Large Language Models (LLMs) is marked by a unique challenge: as these models burgeon in size and complexity, predicting their performance trajectory remains a conundrum. This unpredictability isn’t just a technical hurdle; it accentuates the high resource costs often sunk in trial-and-error endeavors. One potential avenue is to harness the growth trajectory of a rudimentary ‘seed’ model and extrapolate its behavior to larger architectures. However, this method, while resource-conserving, might overlook the non-linearities in performance that larger models exhibit. Alternatively, simulating potential outcomes

from varying scales or making architectural nuances might offer insights, but these simulations are only as reliable as the foundational assumptions they rest upon. Another approach could be to set benchmarking paradigms, rigorously testing model versions across different scales. Such empirical endeavors could hint at underlying scaling laws, but they also come with their hefty computational price tag. Amid these methodologies, the onus is on researchers to judiciously navigate the trade-offs, seeking an optimal interplay between accurate performance predictions and efficient resource utilization.

- **Indistinguishability Between Generated and Human-Written Text** The text produced by large language models (LLMs) can closely resemble human-penned content, leading to concerns about misinformation spread. While it's essential to detect such LLM-generated content to mitigate the risks of misinformation, plagiarism, and impersonation, the improved fluency of these models complicates the detection task. Current countermeasures include post-hoc detectors, which identify unusual tokens to distinguish generated from genuine content, and watermark schemes that modify the text generation process to embed detectable patterns. However, these methods aren't foolproof. For example, adversaries can rewrite LLM-generated content to eliminate unique characteristics, making detection challenging. Techniques like using models to produce synonymous content can aid in this, allowing adversaries to retain the core meaning while altering the structure. Additionally, there's the potential to misuse the watermark schemes: by querying a watermarked LLM repeatedly, one can misclassify authentic human content as LLM-generated, further complicating detection. In summary, while several methods are proposed for detecting machine-generated text, practical application presents numerous challenges, and further research is needed to address these issues.

- **Safety Alignment.** Safety alignment is pivotal for the advancement of AI technologies, particularly for Large Language Models. It goes beyond merely addressing the existential risks of AI. The safety of AI systems, including LLMs, should be woven into their development and deployment. Interpretability, governance, and model property verification are vital components in this process.

LLMs, despite being trained to predict the next word in textual corpora, can infer and represent active attributes of the text authors, like intentions, beliefs, or goals. This brings forward the hypothesis that LLMs can simulate human communicative intents, beliefs, and desires. If true, this accentuates the alignment challenge and potentially introduces new hurdles. There's the looming risk of the models inheriting flawed beliefs, malicious intents, or even pursuing objectives misaligned with human values. Further research is imperative to detect and mitigate such behaviors, ensuring the safe application of LLMs.

In sum, aligning LLMs with human values, goals, and expectations remains a critical challenge. Intensified research efforts are needed, especially in detecting misaligned behaviors and developing strategies to rectify them.

- **Hallucinations** Large language models (LLMs), like ChatGPT, have increasingly been used for day-to-day queries due to their rising popularity. Yet, their accuracy is pivotal as they often produce "hallucinations" or generate inaccurate information. This can be especially deceptive as the fluency of the output can mask these errors. Two types of hallucinations can be identified: *inherent hallucinations*, where the generated text logically contradicts the input, and *external hallucinations*, where the given input doesn't suffice to verify the output's accuracy. Although external hallucinations aren't necessarily incorrect, they are still problematic because of their unverifiable nature. A noticeable challenge in LLMs is the trade-off between diversity and quality. Traditional decoding often introduces randomness at each step, leading to hallucinations. The more diverse the generated answers, the higher the risk of producing hallucinations. Addressing these challenges, some methods, like the Uncertainty-Aware Beam Search, incorporate penalties during decoding for high predictive uncertainties. The Confident Decoding approach suggests

that hallucinations arise from decoding without proper attention to the input. It employs an attention-based confidence score to gauge the model’s focus on input, ensuring high-confidence output generation through a variational Bayesian training process.

- **Limited Context Length** In the realm of language models, especially the large ones, a pertinent issue is their limited context length, hindering their ability to understand and generate extensive texts. This becomes paramount in natural language processing tasks. For example, understanding novels or academic texts isn’t just about a few words or sentences; one needs to consider the whole content. Similarly, deciphering a comment in a meeting transcript might see its meaning swing between sarcasm and seriousness based on preceding discussions. Interestingly, a myriad of open-source models claim robust performance over extended contexts, but in reality, they often falter. Addressing context limitations has brought forth diverse strategies. There’s a push towards devising efficient attention mechanisms, such as linear nested attention [? ], Transient Global attention [? ], CoLT5 [? ], and Synthesizer [? ]. Then, we have position embeddings like the absolute and relative types, including RoPE [? ] and relative positional biases, which look promising but come with their own set of challenges. Alternately, there’s also a discourse around alternative architectures to the Transformer, with state-space models, convolutions, and recurrent neural networks emerging as contenders. They combine computational efficiencies with commendable performance. Tackling these challenges head-on can reshape the landscape of language models, making them more adept for varied applications.
- **Inference Latency** Large language models (LLMs) often experience increased latency due to their intricate architectures. This latency poses a challenge for real-time applications as LLMs may take longer to respond. Primarily, this delay arises from the model’s serial token processing approach and extensive memory consumption, both because of its sheer size and the temporary states, like attention vectors, maintained during decoding. The quadratic scalability of the attention mechanism in Transformers exacerbates this. However, the research community has been proactive in addressing these bottlenecks. Techniques to cut down on memory demands, both in terms of size and bandwidth, and to speed up specific computations have been devised. For instance, some methods modify the attention mechanism to be more hardware-aware or employ higher-order approximations. Quantization stands out as a technique that decreases memory usage and boosts throughput by reducing the computational precision of weights and activations post-training. Pruning, another post-training approach, strategically removes certain model weights without hampering performance. Hybrid expert architectures utilize a collection of specialized modules combined with a routing network to trim down inference time. Cascading leverages models of varying sizes to strike a balance between accuracy and computational overhead. Moreover, the choice of decoding strategy plays a pivotal role in influencing computational costs. To aid in efficient deployment, several frameworks and libraries have been developed that either streamline LLM implementation, curb memory prerequisites, or harness distributed computation strategies.
- **Embodied LLMs.** While LLMs excel at textual and coding tasks, seamlessly integrating them with diverse modalities like vision, speech, knowledge bases, robotic actions, and more remains an uphill battle. Truly embodying AI using LLMs is a stepping stone to achieving artificial general intelligence. Presently, the two prevalent strategies are employing LLMs as modality interfaces or forging ahead with end-to-end multi-modality models. However, both approaches often falter in real-world scenarios due to their inherent limitations. A fusion of these methods, refined data alignment techniques, or novel training paradigms may pave the way for advancements in this domain.
- **Planning and Reasoning.** While LLMs exhibit emergent reasoning capabilities, they lag significantly behind human-like reasoning. Indeed, discerning whether LLMs truly “reason”—by standards of planning, knowledge decomposition, and composition—is a point of contention. Historically, the NLP community tackled reasoning via two main avenues:

incorporating inductive biases into models, exemplified by modular constructs or intermediate hidden abstractions, and decoupling reasoning from modeling, often achieved by creating executable interfaces.

In the context of LLMs, these approaches are mirrored in advanced prompting techniques, such as Chain-of-Thought prompting, which act as inductive biases. Furthermore, leveraging LLMs to create interfaces that integrate external reasoning tools remains a potent strategy. However, the challenge of bestowing LLMs with authentic human-like reasoning remains unresolved. The intricacies of human reasoning itself are nebulous; for instance, the discrete nature of our neuronal processes may hint at our discrete reasoning capabilities. The “all or none” phenomenon of human neuron action potentials—whereby a neuron either fully activates or doesn’t—and the dynamic formation and dissolution of synapses between neurons might offer insights. For neural models to approach human-level reasoning, profound alterations, potentially in activation functions or backpropagation mechanisms, might be imperative.

- **Interactive and Human-centered LLMs.** Large Language Models (LLMs) have the potential to revolutionize many domains, but their practical application isn’t without challenges. A paramount concern is alignment—ensuring LLMs resonate with human values and priorities. This isn’t just about avoiding rogue behaviors, but about truly integrating these models into our workflows in ways that feel natural and beneficial. Indeed, methods ensuring that LLMs behave as intended are crucial, especially as these models become increasingly autonomous. It’s imperative not to let them inadvertently optimize for undesirable outcomes. This calls for a proactive approach: rather than retrospectively fixing misaligned models, alignment techniques should be integral from the onset of model development. Transparency and interpretability of LLMs play pivotal roles in this alignment process. Without a clear understanding of how these models arrive at their conclusions, ensuring their alignment with human values becomes an uphill battle. Looking ahead, the horizon presents an even greater challenge: aligning LLMs that surpass human capabilities. These ‘superhuman’ systems might introduce unprecedented complexities and ethical concerns. As cited in works like [7] and [16], the potential implications of these systems necessitate careful consideration. Furthermore, enhancing the interactivity between humans and LLMs is essential. An effective human-in-the-loop mechanism can be pivotal in maximizing the real-world value of LLMs. It’s worth noting that the bridge between Natural Language Processing (NLP) and Human-Computer Interaction (HCI) needs strengthening. Encouraging collaboration between these communities can drive forward more intuitive and beneficial LLM-human interactions.

## 7 CONCLUSION

Recent advances in large language models have been revolutionizing the field of natural language processing. Effectively using LLMs requires understanding their capabilities, and limitations for various NLP tasks. This work presents a practical guide to working with LLMs for downstream NLP tasks. We first discuss prominent models like GPT-style and BERT-style architectures and the factors influencing their performance. We then explore using LLMs for downstream tasks, including knowledge-intensive tasks, NLU, and NLG tasks, as well as providing concrete examples of successes and limitations. Finally, we indicate the existing challenges, which should be carefully considered in practice. This practical guide offers insights into LLMs and best practices for harnessing LLMs across NLP tasks. We hope it would enable researchers and practitioners to leverage their potential and drive innovation in language technologies.

## ACKNOWLEDGEMENTS

## REFERENCES

- [1] Chatgpt is banned in italy over privacy concerns - the new york times. <https://www.nytimes.com/2023/03/31/technology/chatgpt-italy-ban.html>. (Accessed on 04/23/2023).

- [2] Openai’s gpt-3 language model: A technical overview. <https://lambdalabs.com/blog/demystifying-gpt-3#1>. (Accessed on 03/02/2023).
- [3] Pricing. <https://openai.com/pricing>. (Accessed on 03/02/2023).
- [4] Ahmed Alajrami and Nikolaos Aletras. How does the pre-training objective affect what large language models learn about linguistic properties? In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 131–147, 2022.
- [5] Anil Ananthaswamy. In ai, is bigger always better? *Nature*, 615(7951):202–205, 2023.
- [6] Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- [7] Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022.
- [8] Hangbo Bao, Li Dong, Furu Wei, Wenhui Wang, Nan Yang, Xiaodong Liu, Yu Wang, Jianfeng Gao, Songhao Piao, Ming Zhou, et al. Unilmv2: Pseudo-masked language models for unified language model pre-training. In *International conference on machine learning*, pages 642–652. PMLR, 2020.
- [9] Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. Semantic parsing on freebase from question-answer pairs. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1533–1544, 2013.
- [10] Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Michal Podstawski, Hubert Niewiadomski, Piotr Nyczyk, et al. Graph of thoughts: Solving elaborate problems with large language models. *arXiv preprint arXiv:2308.09687*, 2023.
- [11] Camiel J Beukeboom and Christian Burgers. How stereotypes are shared through language: a review and introduction of the social categories and stereotypes communication (scsc) framework. *Review of Communication Research*, 7:1–37, 2019.
- [12] Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurélie Névél, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos Zampieri. Findings of the 2016 conference on machine translation. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 131–198, Berlin, Germany, August 2016. Association for Computational Linguistics.
- [13] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- [14] Daniel Borkan, Lucas Dixon, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. Nuanced metrics for measuring unintended bias with real data for text classification. In *Companion proceedings of the 2019 world wide web conference*, pages 491–500, 2019.
- [15] Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*, 2015.
- [16] Samuel R Bowman, Jeeyoon Hyun, Ethan Perez, Edwin Chen, Craig Pettit, Scott Heiner, Kamile Lukosuite, Amanda Askell, Andy Jones, Anna Chen, et al. Measuring progress on scalable oversight for large language models. *arXiv preprint arXiv:2211.03540*, 2022.
- [17] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [18] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91. PMLR, 2018.
- [19] Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Kaijie Zhu, Hao Chen, Linyi Yang, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. A survey on evaluation of large language models. *arXiv preprint arXiv:2307.03109*, 2023.
- [20] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- [21] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794*, 2022.
- [22] Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wen-tau Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. Quac: Question answering in context. *arXiv preprint arXiv:1808.07036*, 2018.
- [23] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
- [24] Zheng Chu, Jingchang Chen, Qianglong Chen, Weijiang Yu, Tao He, Haotian Wang, Weihua Peng, Ming Liu, Bing Qin, and Ting Liu. A survey of chain of thought reasoning: Advances, frontiers and future. *arXiv preprint arXiv:2309.15402*, 2023.
- [25] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022.
- [26] Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- [27] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- [28] Haixing Dai, Zhengliang Liu, Wenxiong Liao, Xiaoke Huang, Zihao Wu, Lin Zhao, Wei Liu, Ninghao Liu, Sheng Li, Dajiang Zhu, et al. Chataug: Leveraging chatgpt for text data augmentation. *arXiv preprint arXiv:2302.13007*, 2023.

- [29] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. [arXiv preprint arXiv:1810.04805](#), 2018.
- [30] Bosheng Ding, Chengwei Qin, Linlin Liu, Lidong Bing, Shafiq Joty, and Boyang Li. Is gpt-3 a good data annotator? [arXiv preprint arXiv:2212.10450](#), 2022.
- [31] Jesse Dodge, Taylor Prewitt, Remi Tachet des Combes, Erika Odmark, Roy Schwartz, Emma Strubell, Alexandra Sasha Luccioni, Noah A Smith, Nicole DeCario, and Will Buchanan. Measuring the carbon intensity of ai in cloud instances. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1877–1894, 2022.
- [32] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. A survey for in-context learning. [arXiv preprint arXiv:2301.00234](#), 2022.
- [33] Mengnan Du, Fengxiang He, Na Zou, Dacheng Tao, and Xia Hu. Shortcut learning of large language models in natural language understanding: A survey. [arXiv preprint arXiv:2208.11857](#), 2022.
- [34] Corentin Duchene, Henri Jamet, Pierre Guillaume, and Reda Dehak. A benchmark for toxic comment classification on civil comments dataset. [arXiv preprint arXiv:2301.11125](#), 2023.
- [35] Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. Gptscore: Evaluate as you desire. [arXiv preprint arXiv:2302.04166](#), 2023.
- [36] Yingqiang Ge, Wenye Hua, Kai Mei, Jianchao Ji, Juntao Tan, Shuyuan Xu, Zelong Li, and Yongfeng Zhang. Openagi: When llm meets domain experts. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- [37] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020.
- [38] Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies. *Transactions of the Association for Computational Linguistics*, 9:346–361, 2021.
- [39] Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. Chatgpt outperforms crowd-workers for text-annotation tasks. [arXiv preprint arXiv:2303.15056](#), 2023.
- [40] Tanya Goyal, Junyi Jessy Li, and Greg Durrett. News summarization and evaluation in the era of gpt-3. [arXiv preprint arXiv:2209.12356](#), 2022.
- [41] Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, et al. Textbooks are all you need. [arXiv preprint arXiv:2306.11644](#), 2023.
- [42] Rahul Gupta, Soham Pal, Aditya Kanade, and Shirish Shevade. Deepfix: Fixing common c language errors by deep learning. In *Proceedings of the aaai conference on artificial intelligence*, volume 31, 2017.
- [43] Xiaochuang Han, Daniel Simig, Todor Mihaylov, Yulia Tsvetkov, Asli Celikyilmaz, and Tianlu Wang. Understanding in-context learning via supportive pretraining data. [arXiv preprint arXiv:2306.15091](#), 2023.
- [44] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. [arXiv preprint arXiv:2009.03300](#), 2020.
- [45] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. [arXiv preprint arXiv:2203.15556](#), 2022.
- [46] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. [arXiv preprint arXiv:2106.09685](#), 2021.
- [47] Hang Hua, Xingjian Li, Dejing Dou, Cheng-Zhong Xu, and Jiebo Luo. Fine-tuning pre-trained language models with noise stability regularization. [arXiv preprint arXiv:2206.05658](#), 2022.
- [48] Jie Huang and Kevin Chen-Chuan Chang. Towards reasoning in large language models: A survey. [arXiv preprint arXiv:2212.10403](#), 2022.
- [49] Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. Few-shot Learning with Retrieval Augmented Language Models. 2022.
- [50] Wenxiang Jiao and WenxuanWang Jen-tseHuang XingWang ZhaopengTu. Is chatgpt a good translator? yes with gpt-4 as the engine.
- [51] Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. [arXiv preprint arXiv:1705.03551](#), 2017.
- [52] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. [arXiv preprint arXiv:2001.08361](#), 2020.
- [53] Akhil Kedia, Mohd Abbas Zaidi, and Haejun Lee. Fie: Building a global probability space by leveraging early fusion in encoder for open-domain question answering. [arXiv preprint arXiv:2211.10147](#), 2022.
- [54] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- [55] Tom Kocmi and Christian Federmann. Large language models are state-of-the-art evaluators of translation quality. [arXiv preprint arXiv:2302.14520](#), 2023.
- [56] Lingkai Kong, Haoming Jiang, Yuchen Zhuang, Jie Lyu, Tuo Zhao, and Chao Zhang. Calibrated language model fine-tuning for in-and out-of-distribution data. [arXiv preprint arXiv:2010.11506](#), 2020.
- [57] Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational*

- Linguistics*, 7:453–466, 2019.
- [58] Yuxuan Lai, Chen Zhang, Yansong Feng, Quzhe Huang, and Dongyan Zhao. Why machine reading comprehension models learn shortcuts? In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 989–1002, 2021.
- [59] Hung-Yi Lee, Shang-Wen Li, and Thang Vu. Meta learning for natural language processing: A survey. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 666–684, 2022.
- [60] Yuanyuan Lei and Ruihong Huang. Few-shot (dis) agreement identification in online discussions with regularized and augmented meta-learning. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5581–5593, 2022.
- [61] Yuanyuan Lei, Ruihong Huang, Lu Wang, and Nick Beauchamp. Sentence-level media bias analysis informed by discourse structures. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10040–10050, 2022.
- [62] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*, 2019.
- [63] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021.
- [64] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110*, 2022.
- [65] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.
- [66] Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. Program induction by rationale generation: Learning to solve and explain algebraic word problems. *arXiv preprint arXiv:1705.04146*, 2017.
- [67] Xiao Liu, Kaixuan Ji, Yicheng Fu, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks. *arXiv preprint arXiv:2110.07602*, 2021.
- [68] Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 61–68, 2022.
- [69] Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. Gpteval: Nlg evaluation using gpt-4 with better human alignment, 2023.
- [70] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [71] Yixin Liu, Pengfei Liu, Dragomir Radev, and Graham Neubig. Brio: Bringing order to abstractive summarization. *arXiv preprint arXiv:2203.16804*, 2022.
- [72] Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V Le, Barret Zoph, Jason Wei, et al. The flan collection: Designing data and methods for effective instruction tuning. *arXiv preprint arXiv:2301.13688*, 2023.
- [73] Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. *arXiv preprint arXiv:2104.08786*, 2021.
- [74] Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8086–8098, 2022.
- [75] Andrew Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, pages 142–150, 2011.
- [76] Ian McKenzie, Alexander Lyzhov, Alicia Parrish, Ameya Prabhu, Aaron Mueller, Najoung Kim, Sam Bowman, and Ethan Perez. Inverse scaling prize: Second round winners, 2023.
- [77] Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, Bing Xiang, et al. Abstractive text summarization using sequence-to-sequence rnns and beyond. *arXiv preprint arXiv:1602.06023*, 2016.
- [78] Shashi Narayan, Shay B Cohen, and Mirella Lapata. Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. *arXiv preprint arXiv:1808.08745*, 2018.
- [79] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. Ms marco: A human generated machine reading comprehension dataset. *choice*, 2640:660, 2016.
- [80] Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. Adversarial nli: A new benchmark for natural language understanding. *arXiv preprint arXiv:1910.14599*, 2019.
- [81] OpenAI. Gpt-4 system card.
- [82] OpenAI. Gpt-4 technical report, 2023.
- [83] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- [84] Ankit Pal. Promptify: Structured output from llms. <https://github.com/prompts-lab/Promptify>, 2022. Prompt-Engineering components for NLP tasks in Python.

- [85] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002.
- [86] Arkil Patel, Satwik Bhattamishra, and Navin Goyal. Are nlp models really able to solve simple math word problems? *arXiv preprint arXiv:2103.07191*, 2021.
- [87] Chengwei Qin, Aston Zhang, Zhuosheng Zhang, Jiaao Chen, Michihiro Yasunaga, and Diyi Yang. Is chatgpt a general-purpose natural language processing task solver? *arXiv preprint arXiv:2302.06476*, 2023.
- [88] Shilin Qiu, Qihe Liu, Shijie Zhou, and Wen Huang. Adversarial attack and defense technologies in natural language processing: A survey. *Neurocomputing*, 492:278–307, 2022.
- [89] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.
- [90] Pranav Rajpurkar, Robin Jia, and Percy Liang. Know what you don’t know: Unanswerable questions for squad. *arXiv preprint arXiv:1806.03822*, 2018.
- [91] Siva Reddy, Danqi Chen, and Christopher D Manning. Coqa: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266, 2019.
- [92] Sebastian Ruder, Matthew Peters, Swabha Swayamdipta, and Thomas Wolf. Transfer learning in natural language processing tutorial. *NAACL HTL 2019*, page 15, 2019.
- [93] Erik F Sang and Fien De Meulder. Introduction to the conll-2003 shared task: Language-independent named entity recognition. *arXiv preprint cs/0306050*, 2003.
- [94] Victor Sanh, Albert Webson, Colin Raffel, Stephen H Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, et al. Multitask prompted training enables zero-shot task generalization. *arXiv preprint arXiv:2110.08207*, 2021.
- [95] Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*, 2022.
- [96] Lingfeng Shen, Aayush Mishra, and Daniel Khashabi. Do pretrained transformers really learn in-context by gradient descent? *arXiv preprint arXiv:2310.08540*, 2023.
- [97] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642, 2013.
- [98] Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*, 2022.
- [99] Ruixiang Tang, Yu-Neng Chuang, and Xia Hu. The science of detecting llm-generated texts. *arXiv preprint arXiv:2303.07205*, 2023.
- [100] Ruixiang Tang, Mengnan Du, Yuening Li, Zirui Liu, Na Zou, and Xia Hu. Mitigating gender bias in captioning systems. In *Proceedings of the Web Conference 2021*, pages 633–645, 2021.
- [101] Ruixiang Tang, Xiaotian Han, Xiaoqian Jiang, and Xia Hu. Does synthetic data generation of llms help clinical text mining? *arXiv preprint arXiv:2303.04360*, 2023.
- [102] Ruixiang Tang, Dehan Kong, Longtao Huang, and Hui Xue. Large language models can be lazy learners: Analyze shortcuts in in-context learning. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4645–4657, Toronto, Canada, July 2023. Association for Computational Linguistics.
- [103] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca), 2023.
- [104] Arsene Fansi Tchango, Rishab Goel, Zhi Wen, Julien Martel, and Joumana Ghosn. Ddxplus: A new dataset for automatic medical diagnosis. *Proceedings of the Neural Information Processing Systems-Track on Datasets and Benchmarks*, 2, 2022.
- [105] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothee Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models. 2023.
- [106] Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. Solving math word problems with process-and outcome-based feedback. *arXiv preprint arXiv:2211.14275*, 2022.
- [107] Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. Superglue: A stickier benchmark for general-purpose language understanding systems. *Advances in neural information processing systems*, 32, 2019.
- [108] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
- [109] Jiaan Wang, Yunlong Liang, Fandong Meng, Haoxiang Shi, Zhixu Li, Jianan Xu, Jianfeng Qu, and Jie Zhou. Is chatgpt a good nlg evaluator? a preliminary study. *arXiv preprint arXiv:2303.04048*, 2023.
- [110] Jindong Wang, Xixu Hu, Wenxin Hou, Hao Chen, Runkai Zheng, Yidong Wang, Linyi Yang, Haojun Huang, Wei Ye, Xiubo Geng, et al. On the robustness of chatgpt: An adversarial and out-of-distribution perspective. *arXiv preprint arXiv:2302.12095*, 2023.

- [111] Thomas Wang, Adam Roberts, Daniel Hesslow, Teven Le Scao, Hyung Won Chung, Iz Beltagy, Julien Launay, and Colin Raffel. What language model architecture and pretraining objective works best for zero-shot generalization? In International Conference on Machine Learning, pages 22964–22984. PMLR, 2022.
- [112] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. arXiv preprint arXiv:2208.10442, 2022.
- [113] Albert Webson and Ellie Pavlick. Do prompt-based models really understand the meaning of their prompts? In Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 2300–2344, 2022.
- [114] Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. arXiv preprint arXiv:2109.01652, 2021.
- [115] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. Emergent abilities of large language models. Transactions on Machine Learning Research, 2022. Survey Certification.
- [116] Jason Wei, Yi Tay, and Quoc V Le. Inverse scaling can become u-shaped. arXiv preprint arXiv:2211.02011, 2022.
- [117] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. arXiv preprint arXiv:2201.11903, 2022.
- [118] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Huggingface’s transformers: State-of-the-art natural language processing, 2020.
- [119] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 7959–7971, 2022.
- [120] Jian Yang, Shuming Ma, Haoyang Huang, Dongdong Zhang, Li Dong, Shaohan Huang, Alexandre Muzio, Saksham Singhal, Hany Hassan, Xia Song, and Furu Wei. Multilingual machine translation systems from Microsoft for WMT21 shared task. In Proceedings of the Sixth Conference on Machine Translation, pages 446–455, Online, November 2021. Association for Computational Linguistics.
- [121] Jingfeng Yang, Aditya Gupta, Shyam Upadhyay, Luheng He, Rahul Goel, and Shachi Paul. Tableformer: Robust transformer modeling for table-text encoding. arXiv preprint arXiv:2203.00274, 2022.
- [122] Jingfeng Yang, Haoming Jiang, Qingyu Yin, Danqing Zhang, Bing Yin, and Diyi Yang. Seqzero: Few-shot compositional semantic parsing with sequential prompts and zero-shot models. arXiv preprint arXiv:2205.07381, 2022.
- [123] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. arXiv preprint arXiv:2305.10601, 2023.
- [124] Wengpeng Yin, Jamaal Hay, and Dan Roth. Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. 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), pages 3914–3923, 2019.
- [125] Kang Min Yoo, Dongju Park, Jaewook Kang, Sang-Woo Lee, and Woomyeong Park. Gpt3mix: Leveraging large-scale language models for text augmentation. arXiv preprint arXiv:2104.08826, 2021.
- [126] Jiayi Yuan, Ruixiang Tang, Xiaoqian Jiang, and Xia Hu. Llm for patient-trial matching: Privacy-aware data augmentation towards better performance and generalizability. arXiv preprint arXiv:2303.16756, 2023.
- [127] Daochen Zha, Zaid Pervaiz Bhat, Kwei-Herng Lai, Fan Yang, Zhimeng Jiang, Shaochen Zhong, and Xia Hu. Data-centric artificial intelligence: A survey. arXiv preprint arXiv:2303.10158, 2023.
- [128] Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In International Conference on Machine Learning, pages 11328–11339. PMLR, 2020.
- [129] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. arXiv preprint arXiv:2205.01068, 2022.
- [130] Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B Hashimoto. Benchmarking large language models for news summarization. arXiv preprint arXiv:2301.13848, 2023.
- [131] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. arXiv preprint arXiv:2303.18223, 2023.
- [132] Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. Calibrate before use: Improving few-shot performance of language models. In International Conference on Machine Learning, pages 12697–12706. PMLR, 2021.
- [133] Qihuang Zhong, Liang Ding, Juhua Liu, Bo Du, and Dacheng Tao. Can chatgpt understand too? a comparative study on chatgpt and fine-tuned bert. arXiv preprint arXiv:2302.10198, 2023.
- [134] Ce Zhou, Qian Li, Chen Li, Jun Yu, Yixin Liu, Guangjing Wang, Kai Zhang, Cheng Ji, Qiben Yan, Lifang He, et al. A comprehensive survey on pretrained foundation models: A history from bert to chatgpt. arXiv preprint arXiv:2302.09419, 2023.
- [135] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2022.