

Challenges of Using Pre-trained Models: the Practitioners' Perspective

Xin Tan

*School of Computer Science and Engineering
Beihang University
State Key Laboratory of Complex & Critical Software Environment
Beijing, China
xintan@buaa.edu.cn*

Taichuan Li*

*School of Computer Science and Engineering
Beihang University
Beijing, China
taichuanli@buaa.edu.cn*

Ruohe Chen*

*School of Economics and Management
Beihang University
Beijing, China
ruohechen@buaa.edu.cn*

Fang Liu[†], Li Zhang

*School of Computer Science and Engineering
Beihang University
State Key Laboratory of Complex & Critical Software Environment
Beijing, China
fangliu@buaa.edu.cn, lily@buaa.edu.cn*

Abstract—The challenges associated with using pre-trained models (PTMs) have not been specifically investigated, which hampers their effective utilization. To address this knowledge gap, we collected and analyzed a dataset of 5,896 PTM-related questions on Stack Overflow. We first analyze the popularity and difficulty trends of PTM-related questions. We find that PTM-related questions are becoming more and more popular over time. However, it is noteworthy that PTM-related questions not only have a lower response rate but also exhibit a longer response time compared to many well-researched topics in software engineering. This observation emphasizes the significant difficulty and complexity associated with the practical application of PTMs. To delve into the specific challenges, we manually annotate 430 PTM-related questions, categorizing them into a hierarchical taxonomy of 42 codes (i.e., leaf nodes) and three categories. This taxonomy encompasses many PTM prominent challenges such as fine-tuning, output understanding, and prompt customization, which reflects the gaps between current techniques and practical needs. We discuss the implications of our study for PTM practitioners, vendors, and educators, and suggest possible directions and solutions for future research.

I. INTRODUCTION

Pre-training is a technique that leverages large amounts of unlabeled data to train a deep neural model, which can then be fine-tuned for specific downstream tasks with less labeled data [1]. Pre-training has achieved remarkable success in various domains, especially in natural language processing (NLP) [2] and software engineering (SE) [3], where models such as BERT [4], GPT [5], and T5 [6] have set new state-of-the-art results on many benchmarks. Especially in recent years, the birth of artificial intelligence (AI) tools such as ChatGPT¹, GitHub Copilot², *etc.*, pushed the PTMs to a climax.

PTMs have seen tremendous success recently due to their capability to simplify complex tasks and improve performance through pre-training and fine-tuning. Take GitHub Copilot, a most widely used AI Coding Assistant tool, as an example: GitHub Copilot is powered by Codex [7], a generative pre-trained AI model created by OpenAI. It has been trained on natural language text and source code from publicly available sources. Developers who use Copilot report that they completed the task 55% faster than the developers who did not use GitHub Copilot [8]. Given the immense potential of PTMs, many researchers and AI practitioners are trying to apply PTMs to their own tasks.

Although PTMs are popular and effective, fully harnessing their potential can be challenging for practitioners. On the one hand, PTMs often have complex architectures and mechanisms, making it difficult for practitioners to gain a deep understanding of their inner workings [9]. On the other hand, predefined architectures and computational resource requirements can hinder practitioners from tailoring PTMs to their specific needs or incorporating domain-specific knowledge effectively [10]. Worse still, PTM's community support is currently lacking in maturity because of its rapid development. These challenges are evident in the frequent questions raised by developers in QA forums regarding the use of PTMs. However, there has yet to be a systematic examination of these challenges, which hinders the effective utilization and improvement of PTMs.

To better understand practitioners' challenges and needs, we conduct a large-scale analysis of Stack Overflow (SO) questions related to pre-training. SO³ is the most popular online platform where programmers can ask and answer questions about various topics in software development. It is a

*These authors contributed equally to this work.

[†]Corresponding author.

¹<https://openai.com/blog/chatgpt>

²<https://github.com/features/copilot>

³<https://stackoverflow.com/>

valuable source of data for researchers who want to study the trends, patterns, and challenges in SE [11]–[14]. Specifically, we collect 5,896 PTM-related questions from SO, covering a time span from 2018 to 2023. We perform various statistical and qualitative analyses on this dataset, aiming to answer the following research questions:

- RQ1 **Popularity Trend.** What is the popularity trend of PTM-related questions?
- RQ2 **Difficulty.** What is the difference in difficulty between PTM-related questions and other questions?
- RQ3 **Taxonomy of Challenges.** What specific challenges do PTM practitioners face? Are these challenges different from traditional deep learning challenges?

To address RQ1, we extract and tally the occurrences of various PTM-related tags in the questions, ranking them according to their frequency of mention. To answer RQ2, we compare the response rate, response time, and ratio of answers to views of PTM-related questions and questions unrelated to PTMs, allowing us to assess the level of difficulty. Our findings reveal a consistent upward trend in the number of PTM-related questions over time. PTM-related questions also generally require significantly more time to answer, indicating a higher level of difficulty. To delve into the specific challenges (RQ3), we perform a thematic analysis [15] of the PTM-related questions, which allows us to extract and categorize the main challenges. We establish a comprehensive taxonomy consisting of three categories and 42 codes, linking challenges with Model Lifecycle Management, Coding-related Issues, and Others. Compared with the challenges faced by deep learning (DL) practitioners, PTM practitioners also face various unique prominent challenges, e.g., fine-tuning, output understanding, and memory management. The contributions are as follows:

- We present the first systematic analysis of PTM-related questions on SO, providing a comprehensive overview of the common challenges and needs of PTM practitioners.
- We provide valuable insights and guidance for PTM practitioners, vendors, and educators, and suggest future research directions and opportunities.
- We create and release a dataset of 5,896 questions related to PTMs, which is annotated with categories and metadata. The dataset can be used for further research.

II. BACKGROUND

A. PTMs

PTMs are deep neural network models that are trained on large amounts of unlabeled data, such as text, images, or speech, to learn general and reusable representations or knowledge [16]. These representations or knowledge can then be transferred to specific downstream tasks with less labeled data, by fine-tuning the PTMs with task-specific layers or objectives [9]. PTMs have witnessed remarkable advancements across different fields, particularly in NLP and SE domains [2]. Models like BERT [4], GPT [17], and T5 [6] have demonstrated exceptional performance, establishing new state-of-the-art benchmarks in these areas.

PTMs are categorized into two types based on pre-training objectives: generative PTMs and discriminative PTMs [9]. Generative PTMs learn the joint probability distribution of input data, such as words or pixels, and generate new data samples that follow the same distribution [18]. They typically use autoregressive or autoencoding objectives to maximize the likelihood of input data [19]. For instance, GPT is a generative PTM used for text data, employing a transformer decoder to predict the next word given the previous words [5]. Discriminative PTMs aim to learn the conditional probability distribution of output labels given the input data, such as sentiment polarity or object category [20]. They usually adopt contrastive or cloze objectives to maximize the likelihood of the output labels [19]. For example, BERT is a discriminative PTM used for text data, utilizing a transformer encoder to predict masked words given the surrounding words [4].

B. Pre-training Toolkits

Pre-training toolkits are essential for PTMs by providing the necessary infrastructure and resources. They offer a comprehensive range of functionalities and interfaces that enable easy implementation and utilization of PTMs. Pre-training toolkits usually support various pre-training objectives, architectures, datasets, and downstream tasks, allowing users to customize their configurations and parameters. Pre-training toolkits also provide optimization techniques and distributed computing mechanisms to improve the efficiency and scalability of PTMs. Some popular pre-training toolkits include HuggingFace Transformers [21], Fairseq [22], *etc.*

III. DATA PREPARATION

We focus on SO, which is the most trusted and popular QA website for programmers and has been an important study subject for many SE-related research [11], [13], [23]. As of March 2022, SO has over 20 million registered users and has received over 24 million questions and 35 million answers [24]. Given its popularity among the developer community, SO is well-suited to analyze the challenges that practitioners face when applying PTMs. In this section, we describe how we collect related questions from SO. To collect the data, we used SO API⁴. Considering the popular time span of PTMs, we established a data collection window spanning from January 1, 2018, to June 15, 2023.

We utilized the snowball sampling technique [25] to gather PTM-related questions, initiating with an initial tag list and expanding it based on relevance. After brainstorming with all authors, who have over three years of experience with PTMs, we compiled a list of initial PTM-related tags. The initial set consisted of 10 tags and most of them are directly the name of popular PTMs (i.e., “bert-language-model”, “gpt-2”, “gpt-3”, “roberta”, “roberta-language-model”, “bart”, “t5-transformer”, “huggingface-transformers”, “huggingface-tokenizers”, “huggingface-datasets”). Through this approach, we collected a total of 4,303 PTM-related

⁴<https://api.stackexchange.com/>

questions. Then, we considered the related tags recommended by the SO platform itself to expand the initial tag set. We discussed each recommended tag together to decide whether it was a PTM-related tag. If a recommended tag was related to PTMs, we added it to our tag set and used it to collect more questions. We iterated this process until there were no PTM-related tags recommended by SO. Eventually, we collected 18 tags, mainly involving three types, i.e., names of certain PTMs, PTM provider – *Hugging Face*, and PTM vendor – *OpenAI*. Table I shows the name and frequency of each tag. In total, we obtained 5,896 PTM-related questions. The manual annotation conducted for RQ3 (refer to Section VI-B) demonstrates a low false positive rate (5.8%), providing evidence for the validity of our tag selection process.

TABLE I
NAMES AND FREQUENCIES OF PTM-RELATED TAGS

Selected Tags	No.(Freq.)	Selected Tags	No.(Freq.)
huggingface-transformers	2,563 (43%)	gpt-2	174 (3%)
bert-language-model	1,763 (30%)	sentence-transformers	142 (2%)
openai-api	1,053 (18%)	fine-tune	114 (2%)
pre-trained-model	423 (7%)	openai-whisper	107 (2%)
huggingface-tokenizers	413 (7%)	roberta-language-model	62 (1%)
huggingface	403 (7%)	roberta	32 (1%)
gpt-3	245 (4%)	gpt-4	22 (0.3%)
chatgpt-api	219 (4%)	bart	14 (0.2%)
huggingface-datasets	187 (3%)	t5-transformer	6 (0.1%)

IV. RQ1: POPULARITY TREND

A. Motivation

Despite the growing popularity of PTMs, we still do not know the fluctuating trends and interests surrounding this field on SO. This lack of understanding hinders our ability to gain valuable insights into the evolution of PTMs.

B. Approach

To analyze the popularity trend of PTMs, we utilized the PTM-related questions collected in Section III. These questions were grouped by year and by tag, enabling us to determine the number of questions for each group. By doing so, we gathered the frequency of each tag within each year. This approach allowed us to gain insights into the evolving interests and trends surrounding PTMs over time.

C. Results

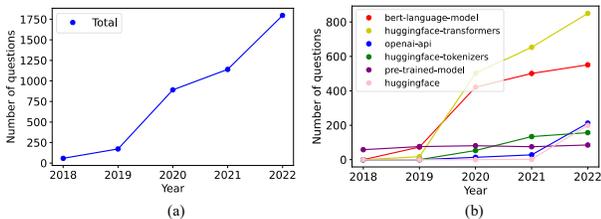


Fig. 1. Popularity Trend of PTM-Related Questions and Top Six Tags

Fig. 1 (a) depicts the popularity trend of PTM-related questions. The graph highlights a consistent upward trend, indicating the increasing popularity of PTMs over the years. Before 2019, the number of PTM-related questions was relatively low, reflecting a limited awareness and adoption of PTMs. However, as time passed by, there is a significant surge in the number of questions, indicating a growing interest and

engagement with PTM technology. Notably, in 2022, there are more than 2,000 questions reported, reflecting the heightened attention and demand for information about PTMs. This substantial increase further reinforces the growing recognition and significance of PTMs in various domains. Fig. 1 (b) depicts the evolution of the six most popular tags, demonstrating a general upward trend. Specifically, the tags “*huggingface-transformers*” and “*bert-language-model*” experienced significant growth from 2019 to 2020. Similarly, the tags “*pre-trained-model*” and “*huggingface*” saw substantial increases from 2021 to 2022.

Summary for RQ1: The number of PTM-related questions has been consistently increasing over the years, indicating the growing recognition and widespread adoption of PTMs, as well as the significant challenges that practitioners face when working with these models.

V. RQ2: DIFFICULTY

A. Motivation

RQ2 examines the difficulty of PTM-related questions in comparison to other topics. It can provide insights into the level of challenge, helping assess the need for additional support to enhance user experience.

B. Approach

In order to comprehend and quantify the difficulties encountered by developers when using PTMs, we employed the following three widely adopted metrics [26]–[28]:

- Percentage of the questions with no accepted answer (i.e., *%no acc*).
- Response time required to receive an acceptable answer.
- The ratio of the average number of answers these questions receive to the average number of views these questions receive (i.e., *PD*) [28].

To establish a baseline for comparison, we utilized questions posted on SO after January 1, 2018, excluding PTM-related questions, which we denoted as “PTM-unrelated questions”. For the first metric, we conducted a proportion test [29] to compare the percentages of questions with no accepted answers between PTM-related and PTM-unrelated questions. This statistical test is appropriate for comparing the *%no acc* values and evaluates the null hypothesis that the proportions are equivalent across multiple groups. Regarding the second metric, we analyzed questions with accepted answers and visualized the distribution and median response time for both PTM-related and PTM-unrelated questions. The third metric involves calculating the average values of “AnswerCount” and “ViewCount” for PTM-related and PTM-unrelated questions to obtain *PD* values. Generally, when a question obtains a high number of views but receives only a limited number of answers, it suggests that only a small fraction of the audience possesses the expertise to respond adequately [28]. Thus, a lower score of *PD* indicates a higher level of difficulty for questions within that topic.

C. Results

For the percentage of questions without an accepted answer (*%no acc*), we find that PTM-related questions have a value of 73.9%, while PTM-unrelated questions have a value of 57.8%. The proportion test results ($\chi^2 = 628.830$, $df = 1$, $p\text{-value} < 1e - 100$) confirm the significant difference, indicating that PTM-related questions are more challenging to answer compared to PTM-unrelated questions. Table II shows *%no acc* of the top six popular tags. We observe that the questions tagged with “*openai-api*” and “*huggingface*” exhibit the highest values of *%no acc*, indicating that they are most difficult to answer. In terms of this metric, PTM-related questions also present greater difficulties than well-researched topics in SE, such as big data (*%no acc* = 60.5%) [30], concurrency (*%no acc* = 43.8%) [26], and mobile devices (*%no acc* = 55.0%) [27].

TABLE II
%no acc AND *PD* OF PTM-RELATED TAGS

Selected Tags	<i>%no acc</i>	<i>PD</i>
huggingface-transformers	0.70	0.05
bert-language-model	0.74	0.05
openai-api	0.78	0.08
pre-trained-model	0.68	0.06
huggingface-tokenizers	0.70	0.05
huggingface	0.78	0.12

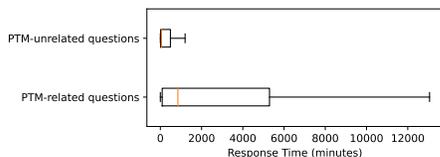


Fig. 2. Time Needed to Receive an Accepted Answer

Fig. 2 illustrates the distribution of response time required to receive accepted answers for PTM-related questions and PTM-unrelated questions. The response time for PTM-unrelated questions predominantly falls below 1,500 mins, while the distribution range is significantly wider for PTM-related questions. We also discover that the average response time for PTM-related questions is more than twice that for PTM-unrelated questions, with 22,676 mins and 9,745 mins, respectively. The median response time for PTM-related questions is 851 mins, which is over 15 times longer than 55 mins for PTM-unrelated questions. In terms of this metric, PTM-related questions also present greater difficulties than well-researched topics in SE, such as big data (198 mins) [30], concurrency (42 mins) [26], and mobile devices (55 mins) [27].

For the third metric, we obtained the *PD* scores of PTM-related and PTM-unrelated questions as 0.06 and 0.12 respectively. Table II shows *PD* of the questions with the top six popular tags. Even when compared to the most challenging subtopics in the field of Security (e.g., “*Java Security*”, “*Asymmetric Encryption*”, and “*Bug*”), where the minimum *PD* value is 0.07 [28], the difficulty level of PTM-related questions with certain tags remains lower. This suggests that answering PTM-related questions poses a greater challenge, requiring a deeper understanding and expertise.

Summary for RQ2: Our findings suggest that developers encounter significant challenges when utilizing PTMs, as indicated by the response rate, response time, and the ratio of answers to views. These challenges highlight the need for further research to better understand the specific difficulties underlying PTM-related questions. Such research can provide valuable insights and pave the way for addressing these challenges effectively.

VI. RQ3: TAXONOMY OF CHALLENGES

A. Motivation

As PTM-related questions gain increasing attention and prove to be difficult to resolve, it is crucial to investigate the particular challenges encountered by PTM practitioners, especially the unique challenges compared to general DL models. This information can reveal the distinct support required.

B. Approach

1) *Collecting Questions for Manual Analysis:* While we identified 5,896 PTM-related questions in Section III, their quality cannot be guaranteed, directly impacting the quality of the constructed taxonomy. To address this, we refined our dataset by excluding PTM-related questions with a score below 1, determined by SO user votes. A question’s score is calculated by subtracting the number of downvotes from the number of upvotes, with a higher score indicating better question quality [31]. Through this process, we filtered out vague, unclear, trivial, or duplicated questions that were unsuitable for our study, resulting in 2,829 high-quality PTM-related questions. Manually analyzing 2,829 questions is time-consuming and can lead to researcher fatigue, risking decreased accuracy and attention to detail in the annotation process. To address this, we adopted two sampling strategies: random sampling and high-score sampling. **Random sampling** aimed to ensure a diverse and representative distribution of topics. **High-score sampling** aimed to complement the questions with high quality and great significance. We combined these two sampling strategies to balance the trade-off between diversity and significance.

Following previous work [13], [23], we conducted a random sampling of 339 questions from 2,829 high-quality PTM-related questions, aiming to achieve a 95% confidence level and a 5% confidence interval [32]. This ensures that our sample is representative and reliable to reflect the characteristics of the PTM-related questions. Considering that random sampling may not cover the most important and high-quality questions, we decided to supplement the questions with a score greater than or equal to nine. This selection results in 116 questions accounting for 4.1% of high-quality PTM-related questions. We opted for a score of nine as the threshold, strategically aiming to concentrate on the most exemplary issues while considering the feasibility of manual analysis. It is worth noting that there were 25 overlapping questions between these two sampling strategies. In total, we obtained 430 questions as our final sample for annotation and analysis. The size of this dataset is comparable to those used in previous studies [13], [14] that also required manual analysis of SO posts.

2) *Conducting Manual Analysis*: To analyze the challenges within the questions, we conducted a thematic analysis [15]. We initially analyzed a subset of 129 questions (30% of the total 430 questions) by carefully reading and re-reading them, including the title, body, code snippets, comments, and answers. We assigned initial codes to reflect the underlying challenges of these questions, using keywords or conducting further investigation for deeper understanding. We also re-viewed answers and used the Google search engine to gather additional information. If a question is not related to PTM, we classify it as a False positive.

We then grouped similar codes into categories and created a hierarchical taxonomy of challenges. This iterative process involved identifying recurring patterns, concepts, and ideas within the questions and organizing them accordingly. The outcome was a comprehensive codebook illustrating various codes related to PTM challenges and their meaning. This analysis was performed by the first four authors together, i.e., two professors, and two undergraduate students, all with over three years of PTM-related experience. This collaborative approach ensured consistency and reliability in the pilot analysis, providing a data guarantee for the subsequent analysis.

Using the developed codebook, the remaining 70% of questions were labeled independently by the second and third authors. In cases where a question did not fit within the existing taxonomy, the first four authors held discussions to establish new codes and enhance the codebook and taxonomy. In the later stage of the coding process, there were no new codes emerged, which means that we had reached data saturation [33]. To assess the reliability of the labeling process, inter-rater agreement was measured during the independent labeling phase. The agreement, measured using Cohen’s Kappa, was found to be 83%, indicating perfect agreement and highlighting the reliability of our coding schema and procedure. Any conflicts or discrepancies in labeling were resolved through discussions and consensus among the first four authors, ensuring consistency and accuracy in the final classification of the questions.

3) *Mapping to DL challenges*: To comprehend the differences in challenges between PTM and general DL, we aligned our taxonomy with the work of Zhao et al. [34]. They investigated the challenges of DL by analyzing SO and obtained 30 challenges. We identified the analogous challenges and annotated them accordingly. There are certain codes that do not align with our taxonomy. These codes primarily pertain to specific applications (e.g., “*Object Detection*”) or specific DL models (e.g., “*CNN Structure*”).

C. Results

Fig. 3 shows the hierarchical taxonomy of challenges extracted from PTM-related questions. The taxonomy is structured into multiple layers with varying numbers across different categories. We identified a total of 42 challenges as leaf nodes within the taxonomy, which is composed of three primary categories. We can see that PTM practitioners also encounter common challenges in DL [34], such as *Model API*

Usage, Input Data Processing, Embedding, and Model Design. However, due to the uniqueness of PTMs, PTM practitioners face many specific and prominent challenges, such as *Fine-tuning, Output Understanding, and Prompt Customization*. We provide further details on these categories by combining the features of PTMs and giving examples.

I. *Model Lifecycle Management (56.79%)*

This is the most extensive and significant category, covering more than 50% of the PTM-related questions. It consists of nine sub-categories related to the complete lifecycle of a PTM, from its construction to testing and deployment.

1) *Model Construction (19.75%)*. Constructing PTMs can pose challenges for many practitioners in the field. Notably, nearly a quarter of questions pertain to this subcategory. It encompasses various crucial aspects that directly impact the construction of PTMs, e.g., *General Design, Model Component Understanding, and Tokenizer*. Each of these topics plays a vital role in shaping the overall construction and functionality of PTMs.

General Design. Designing the appropriate architecture and approach for PTMs can be challenging for developers. One challenge is *Model Design*, which involves making decisions on the neural network architecture, pre-training objectives, regularization techniques, and model size to effectively apply to downstream tasks. For example, a developer wanted to disable or freeze certain layers by modifying the model’s architecture [35]. Another challenge is *Model Selection*, as the wide range of available PTMs requires careful consideration of factors such as performance, domain similarity, computational requirements, available resources, etc. Developers usually have difficulty in selecting suitable models for their tasks, e.g., “*I want to use Spacy’s pre-trained BERT model for text classification but I’m a little confused about cased/uncased models*” [36]. Developers also feel difficulty in *Model Ensemble* due to the need to consider compatibility, selection, and strategy for optimal integration of multiple PTMs. For example, “*I’m trying to create an ensemble with three pre-trained VGG16, InceptionV3, and EfficientNetB0 for a medical image classification task...But when I execute the code, I get this error...*” [37].

Model Component Understanding. Understanding the components of PTMs can be challenging, especially for developers new to working with such models. Developers often have issues related to the model’s internal information, like concepts and properties. For instance, “*I’ve googled but didn’t find a single guide that allowed me to view how a pre-trained torch neural network is designed/coded*” [38]. Worse still, the internal workings of PTM are like a black box to practitioners, making it difficult to directly observe or explain the specific workings of the model. This makes it challenging to understand the functionality and interactions of these components.

Tokenizer. Developers face challenges related to tokenization when breaking down a sequence of text into smaller units called tokens. These challenges arise when applying tokeniza-

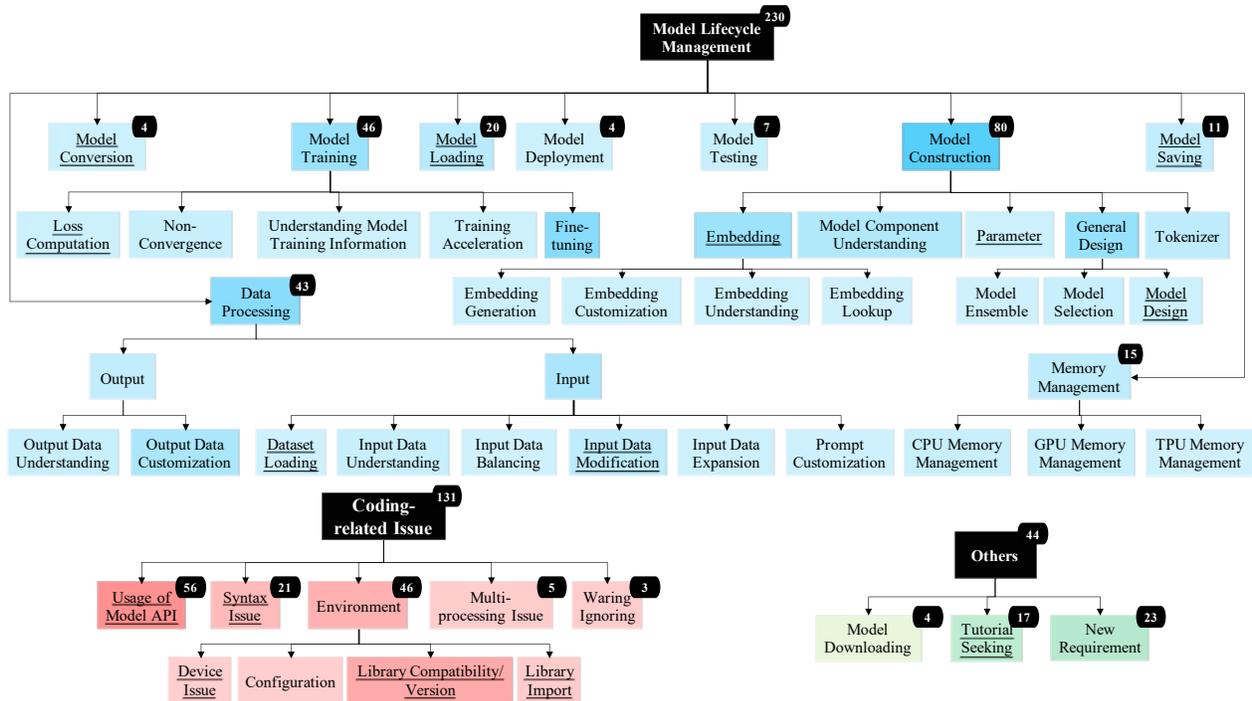


Fig. 3. Taxonomy of PTM Challenges. The challenges highlighted in codes indicate that these challenges are also prevalent in the broader field of deep learning [34], while others are prevalent in PTMs.

tion to a specific task or when encountering errors during the tokenization process. For example, a developer asked “*how do I use ByteLevelBPETokenizer with UTF-8?*” [39]. An example of a common error encountered during tokenization of BERT or DistilBERT models is the *ValueError: "TextEncodeInput must be Union[TextInputSequence, Tuple[InputSequence, InputSequence]]* [40]. his error suggests a problem with the input data format/type provided to the tokenizer. This question has gained significant attention, with 41k views and 35 upvotes. The top-rated answer, receiving 79 upvotes, suggests a solution of *removing all None rows from DataFrame columns before converting them to a list*.

Embedding. An embedding is a numerical vector representation of words, sentences, or documents that captures the semantic and contextual information of the input text [41]. Developers often encounter challenges related to embeddings, including *Embedding Generation*, *Embedding Understanding*, *Embedding Customization*, and *Embedding Lookup*. For instance, a developer had a specific question about generating embeddings using BERT and asked, “*How to cluster similar sentences using BERT?*” [42]. This question has received 34k views and 30 upvotes, indicating a need for examples and guidance on generating embeddings. Developers also face difficulties when customizing their embeddings. For instance, a developer mentioned, “*Earlier I’ve used Glove embedding to build the seq2seq model for text summarization. Now I want to change the Glove with BERT to see the performance of the model?*” [43]. These examples highlight the importance of providing guidance for developers in generating embeddings and customizing them to suit their specific tasks.

Parameter. Only a few developers face challenges when

working with PTM parameters. Developers often need to fine-tune or adapt the PTM parameters to their specific task or dataset. This involves adjusting hyper-parameters, e.g., learning rate, batch size, regularization method, and optimizer parameters like momentum, and learning rate decay. For instance, because some optimizer parameters are not specified correctly, the switch from a custom-built model to a pre-trained BERT model resulted in a significant drop in accuracy [44].

2) Data Processing (10.62%). This subcategory is related to input and output data processing, accounting for over 10% of PTM-related questions, indicating its significance in applying PTMs. This includes challenges such as dataset loading, input data modification, and handling the model’s output to suit the specific needs of the application or downstream tasks.

Input. Developers encounter several challenges when processing input data for PTMs, with *Input Data Modification* being a common obstacle. This involves tasks such as cleaning, preprocessing, formatting, and transforming the data to meet PTMs’ requirements, which can be puzzling for developers. For instance, when using BERT to analyze tweet sentiment, a developer faced the challenge of identifying which elements to remove, e.g., *@names, URLs, or numbers* [45]. Developers also struggle with *Understanding Input Data*, including settings or shapes. For example, a developer questioned the need for a throw-away column in the BERT format [46], highlighting the difficulty in comprehending certain data-related elements or configurations. *Input Data Expansion* presents challenges as well. For example, due to time constraints, a developer sought to incorporate new data in order to continuously train the model instead of starting the training process from scratch [47]. Additionally, a few developers encounter

challenges with *Prompt Customization*, *Dataset Loading*, and *Input Data Balancing*. These challenges emphasize the complexity of working with the input data of PTMs and the diverse aspects that developers need to address.

Output. Developers encounter challenges related to the output of PTMs. These challenges include *Output Data Understanding* and *Output Data Customization*. Output Data Understanding refers to the difficulty in comprehending the format, structure, or interpretation of the model’s output. For instance, a developer asked about the meaning of the second output of Huggingface’s BERT [48]. Compared to DL, PTMs often generate complex output representations that require further analysis or processing to extract meaningful insights. Understanding the output data is crucial for effectively interpreting and utilizing the model’s predictions. Output Data Customization involves the need to customize or adapt the model’s output to fulfill specific requirements or downstream tasks. This process may involve post-processing, format conversion, or integrating additional information. For example, a developer encountered challenges in retrieving all documents related to a topic when using BERTopic for topic modeling, as it only returned three documents per topic [49].

3) Model Training (10.36%). PTM is usually pre-trained on large-scale corpora, but when applied to specific tasks, it needs to be adapted to real-world data through **Fine-tuning**. However, challenges arise in selecting appropriate learning rates, freezing or unfreezing layers, and avoiding overfitting. For example, a developer asked “*How to train BART for text summarization using custom dataset?*” [50]. Sometimes, developers want to fine-tune PTMs to new tasks. This is more challenging, as these new tasks are often cross-field, not just new datasets in the same field. Thus, major modifications usually need to be made to the model to meet the demand. For example, “*I fine-tuned a BERT (or RoBERTa) model for sequence classification. Can I fine-tune the same model for a different task (QA or Sentiment Analysis)?*” [51].

Another significant challenge that developers face during model training is **Non-Convergence**. Non-convergence refers to the situation where the model fails to reach an optimal solution during training. This can manifest in various ways, such as the loss function not decreasing, the model’s performance plateauing, or the training process getting stuck. For instance, a developer encountered non-convergence when applying a pre-trained HuggingFace ALBERT transformer model to a text classification task, as the loss was not decreasing [52].

Compared to DL, training PTMs can be more computationally intensive and time-consuming, making **Training Acceleration** crucial for efficient model development. Developers often struggle with ways to accelerate the training process. For example, a developer asked, “*Is there any way to do some CPU optimization to reduce the training time?*” [53]. Developers also face challenges in **Understanding Model Training Information**, including concepts and methods. This difficulty highlights the need for clarity and comprehension of training-related information. Additionally, challenges of **Loss Computation** also arise within this subcategory.

4) Model Loading (4.93%). The challenge of model loading refers to the difficulties developers may face when loading PTMs into memory. It can be challenging due to factors such as large model sizes, disk I/O bottlenecks, and path issues [54]. Developers can address this challenge by considering techniques such as model compression, incremental loading, and model caching. These strategies help optimize the loading process and enhance the integration of PTMs into applications.

5) Memory Management (3.70%). Developers need to manage memory when working with PTMs to ensure efficient loading, execution, and training. This is crucial due to the large sizes of PTMs, memory requirements during inference and fine-tuning, as well as the demands of batch processing and parallel processing. However, developers often encounter challenges related to memory management, including **CPU, GPU, or TPU management**. These challenges arise from various factors such as insufficient memory capacity, memory allocation, model parallelism, batch processing, and TPU memory utilization. For instance, a developer asked “*How can I make full utilization of RAM & GPU memory?*” [55]. Developers can optimize memory management by exploring techniques like data parallelism, memory optimization algorithms, and efficient batch processing. They can leverage memory-efficient tools, frameworks, and hardware-specific optimizations. Effective memory management helps overcome challenges and maximizes resource utilization for PTMs.

6) Other Challenges related to Model Lifecycle Management (6.42%). Besides the challenges discussed earlier, developers may also encounter issues related to **Model Saving**. Saving a trained model is crucial for future use, deployment, or sharing with others. However, challenges may arise when saving the model weights with customized requirements. These requirements can include saving and averaging weights of models [56], or saving only the weights that achieve the best validation performance [57]. **Model Testing** is essential to ensure high performance, involving metrics selection, dataset preparation, and comprehensive analysis. Developers may feel puzzled about Model Testing. A developer asked, “*how to do a test before fine-tuning, in order to compare two models before and after fine-tuning?*” [58]. Other challenges include **Model Deployment** and **Model Conversion**. Model deployment integrates trained models into a production environment, while Model Conversion converts models between libraries/formats. For instance, developers aim to improve Hugging Face transformer model inference speed without GPU [59] or inquire about converting a BERT model to TFLite [60].

II. Coding-related Issues (32.35%)

When developers utilize PTMs, they may come across various coding-related challenges. These challenges encompass a range of issues, such as correctly using the model’s API, ensuring compatibility with the coding environment, handling syntax errors, addressing resource-intensive processing demands, and appropriately addressing warnings. Developers can effectively navigate the complexities of using PTMs by being mindful of potential issues and optimizing their code.

1) **Usage of Model API (13.83%)**. Developers often encounter challenges related to understanding, generating, or modifying code implementation with PTM APIs. This is a widespread issue that impacts over 10% of PTM-related questions. One key difficulty lies in the complexity of PTM APIs, which necessitates a thorough comprehension of their documentation and effective utilization within code. Interacting with a PTM API involves understanding its functionalities, parameters, and input-output data formats. It is essential to have a solid grasp of the underlying structure and functionality of the PTM to effectively work with the API. We find that most of the questions are related to parameters. For instance, a developer who followed a basic tutorial on the ChatGPT API encountered the following issue: “*OpenAI ChatGPT (GPT-3.5) API error: ‘InvalidRequestError: Unrecognized request argument supplied: messages’*” [61]. This error indicates that the API request included an unrecognized argument, possibly related to the “messages” parameter. Another common issue stems from developers being unclear about the type of parameter expected by an API. For example, a developer encountered the error “*TypeError: not a string — parameters in AutoTokenizer.from_pretrained()*” while trying to use the *AutoTokenizer.from_pretrained()* function [62]. Additionally, developers may struggle to understand how to use a specific API to implement a particular task. For instance, a developer expressed uncertainty about which API to use when attempting to implement the OpenAI GPT-3 API Client in PHP [63]. The above challenges indicate that the current API documentation may sometimes be insufficient.

2) **Environment (11.36%)**. Developers may face environment-related issues when working with PTMs, such as compatibility problems with different versions of programming languages, frameworks, or libraries. These issues can arise due to differences in dependencies, package versions, or even operating systems.

Library Compatibility/Version. When working with PTMs, many questions about the Environment often stem from compatibility or version issues with libraries. Incompatibilities between the frameworks or libraries and the PTMs’ requirements can cause various problems. Here are some common scenarios. PTMs typically need specific libraries or frameworks to be installed, and using an incompatible version can lead to errors or unexpected behavior. For example, a developer encountered an error (*openai error: subprocess-exited-with-error*) while attempting to install openai with Python 3.11 on Windows OS [64]. The root cause in this case could be the need for a newer version of pip to ensure the proper installation of openai. Upgrading a library version may also introduce changes to the corresponding API. If developers are not promptly informed of these changes, errors may occur. For instance, when a developer tried to fine-tune a Huggingface GPT2LMHeadModel model for casual language modeling using PyTorch Lightning, they encountered the error: *AttributeError: ‘str’ object has no attribute ‘size’* [65]. The solution to this issue would involve using a recent version of the “transformer” library, as it likely introduced changes to the

API that were not accounted for in the code.

Developers may also face the challenges with **Library Import**. It means that they do not know how to properly import a certain library. The possible reasons include unfamiliarity with library names, incorrect import syntax, incorrect installation methods, or importing outdated libraries. For instance, a developer may face an issue if they fail to install the “transformer” library using its official Git link. This can result in errors when attempting to use certain functions from the library [66]. In addition to the challenges mentioned above, developers may also face environmental issues in two other areas. **Configuration**: Configuration-related challenges involve difficulties in setting up or configuring the PTM for specific tasks or use cases. Developers may need to refer to the PTM’s documentation or seek guidance from the community to overcome configuration hurdles. **Device Issue**: When working with PTMs, device-related challenges can arise due to hardware limitations or compatibility issues. These challenges may include memory constraints, GPU availability, GPU driver version, and optimizing the code for efficient usage of resources. To address these issues, developers may need to adjust batch sizes, utilize distributed computing, or explore hardware-specific optimizations.

3) **Other Challenges related to Coding-related Issues (7.16%)**. When using PTMs, developers face coding-related challenges beyond the Environment and Usage of Model API. These challenges include **Syntax Issue**, **Multi-processing Issue**, and **Warning Ignoring**. Syntax issues result from errors in code structure or format. These issues are often caused by a lack of programming experience or carelessness. Multi-processing challenges arise when developers need to manage the concurrent or parallel execution of code involving PTMs. This can occur in situations where multiple requests must be processed simultaneously or when utilizing distributed computing for quicker inference. For example, a developer used Hugging Face’s zero-shot classification pipeline for question answering but found that implementing multiprocessing did not improve the timing compared to using a for loop [67]. Moreover, PTM-related libraries often provide warning messages to alert developers about potential issues or best practices. Sometimes, developers want to ignore these warnings.

III. Others (10.86%)

The third category of PTM-related questions encompasses three subcategories of issues and challenges that typically fall outside the scope of the specific challenges mentioned earlier.

We observe that 5.68% of developers propose **New Requirement** for PTMs. This challenge occurs when developers have specific use cases or requirements that are not currently supported by existing PTM-related APIs. In such cases, developers are interested in understanding how to implement these desired functionalities or features using PTMs. For instance, developers may seek to customize positional embeddings [68] or integrate PTMs with specific data formats like JSON [69], encountering challenges along the way. Another subcategory of challenges, accounting for 4.20%, is **Tutorial Seeking**.

Developers often face situations where they need to complete a specific task but lack the knowledge or guidance on how to do so. In such cases, they seek tutorials or guidance from the community to overcome these challenges and accomplish their objectives. For example, a developer asked, “*I would like to train an encoder-decoder model as configured below for a translation task. Could someone guide me as to how I can set up a training pipeline for such a model?*” [70]. Furthermore, a small percentage of developers (1.23%) encounter challenges related to **Model Downloading**. These challenges may involve difficulties in downloading or accessing specific PTM models, which can hinder the progress and usage of the models.

Summary for RQ3: We construct a comprehensive taxonomy of PTM-related challenges, which can be grouped into three categories. The richest category is Model Lifecycle Management, which includes various aspects from model construction to deployment, with nine sub-categories. We also observe that a notable number of developers face coding issues and other challenges such as model downloading. Many challenges are often unique or more serious than those encountered by general deep learning practitioners, such as “*fine-tuning*”, “*memory management*”, and “*model component understanding*”.

VII. IMPLICATIONS

A. Implications for Researchers

We conduct a comprehensive and systematic analysis of the challenges of PTM-related questions on SO, which can help researchers identify current limitations in existing PTMs, as well as inspire new directions and solutions for future research. For example, our analysis reveals that many practitioners struggle with understanding the internal workings of PTMs, such as embeddings, weights, and outputs. This suggests that **more interpretable and explainable PTMs** should be explored, as they can offer greater transparency and insights for users. Additionally, our analysis shows that many practitioners face difficulties customizing or adapting PTMs for new tasks or scenarios, such as changing input or output data formats, adding new training data, or fine-tuning model parameters. This highlights the need for **more flexible and adaptable PTMs** that can accommodate a wider range of applications and use cases. We also found that incompatibilities between frameworks, libraries, and PTM requirements can pose challenges and hinder the smooth integration and utilization of PTMs. To address this issue, future research could focus on developing **automated tools to detect and mitigate library conflicts**, as well as providing mechanisms to resolve conflicts and suggest alternative libraries. Moreover, although SO is a representative QA platform for PTM practitioners, future research can **supplement our analysis with data from additional platforms and repositories** to obtain a more comprehensive understanding of the PTM domain.

B. Implications for Practitioners

Our taxonomy is a helpful guide for developers who are working with PTMs. They can use the taxonomy combined

with our annotated questions to anticipate and overcome specific challenges that they may face during the development and deployment of PTMs. For instance, to tackle challenges in **model construction**, developers can get guidance on model design, pre-training goals, and selecting suitable PTMs for particular tasks. They can also learn techniques for **model customization and generalization** to new tasks. In data processing, developers can refer to resources for modifying input data, understanding input data formats, and expanding input data. They can also learn techniques for customizing the model’s output and handling output data for downstream tasks. In case of challenges in **model training**, developers can pay close attention to non-convergence problems and training acceleration on the basis of understanding model training information. They can also learn techniques for managing model loading and optimizing memory usage. In addition to the above, developers can share their experiences, solutions, and insights with the community by **participating in online forums and engaging with PTM vendors**. They can contribute to the collective knowledge and help others overcome similar challenges. Additionally, our study provides a quantitative analysis of the popularity and difficulty trends of PTM-related questions on SO. This analysis can help practitioners **stay informed about the latest developments and best practices of PTMs and techniques in the field**.

C. Implications for PTM Vendors

The taxonomy of challenges offers valuable insights for PTM vendors to **improve their offerings and provide better support to developers**. PTM vendors can enhance their products by addressing challenges, such as model loading, memory management, and compatibility with different frameworks and libraries. They can optimize model loading processes, develop memory-efficient algorithms, and provide compatibility guidelines and tools. Additionally, they can **create comprehensive documentation, tutorials, and resources** to support developers with challenges related to model construction, data processing, and training. By offering clear explanations, code examples, and troubleshooting guides, vendors can help developers work effectively with their PTM offerings. **Engaging with the developer community** is crucial for PTM vendors. They should establish support channels and bug-reporting mechanisms to listen to developers’ needs and challenges. By incorporating feedback and promptly addressing issues, vendors can build strong relationships with developers and foster a collaborative environment for PTM development.

D. Implications for Educators

The analysis of challenges related to PTMs on SO has important implications for PTM educators, such as instructors and online course providers. It offers valuable insights into the specific areas where learners may encounter difficulties and can guide the design and delivery of educational materials and courses. Specifically, PTM educators can **tailor their curricula** to address the challenges identified in the taxonomy. For example, they can develop instructional materials that focus on explaining the internal workings of PTMs in a

more interpretable and explainable manner. This can include providing visualizations, intuitive explanations, and real-world examples to enhance learners’ understanding of PTM mechanisms like embeddings, weights, and outputs. Furthermore, educators can **design practical exercises and assignments** that allow learners to practice customizing and adapting PTMs for different tasks and scenarios. By incorporating hands-on activities, learners can gain experience in modifying input/output data formats, fine-tuning models, and addressing compatibility issues with libraries and frameworks. Additionally, educators can **provide resources and guidance on library conflict detection and mitigation**. This can involve teaching learners techniques to identify and resolve conflicts, as well as highlighting best practices for managing dependencies and ensuring compatibility between PTMs and various libraries. By incorporating these insights into their educational materials and courses, PTM educators can better prepare learners to overcome the challenges they may encounter when working with PTMs. This will contribute to a more effective and efficient learning experience, enabling learners to confidently apply PTMs in their own projects and research.

VIII. RELATED WORK

A. Studies on SO Data

Many studies use SO data to analyze different aspects of software development, such as programming languages [71], software testing [72], software security [73], and software privacy [74]. Some studies also focus on analyzing specific types of software, such as DL software [75] and mobile applications [27], [76]. Some studies investigate the quality of questions and answers on SO to help developers ask high-quality questions and provide good answers [77]. Some studies also treat SO as a social platform to explore the social activities of developers [78]. These studies provide valuable references for the design of our study.

B. Studies on PTMs

Many studies aim to understand, improve, and apply PTMs. Some studies focus on understanding the behavior and reasoning of PTMs, such as probing their linguistic knowledge [79], analyzing their attention patterns [80], and visualizing their hidden states [80]. Some studies have focused on improving the performance and efficiency of PTMs, such as designing better architectures [81], optimizing better objectives [82], and reducing computation and memory costs [83]. Some studies have focused on applying PTMs to various downstream tasks, such as natural language understanding [84], natural language generation [85], and question answering [86].

Our work complements these studies by providing a practitioner’s perspective on the challenges of using PTMs. We collect and analyze real-world questions posted by developers who use PTMs for various purposes. We identify the common problems and difficulties that they encounter when using PTMs, such as selecting the appropriate PTMs, fine-tuning it for specific tasks, deploying it to different platforms, etc.

IX. THREATS TO VALIDITY

Selection of data source. Our study refers to the previous work [12]–[14], [87] to use SO as our data source to extract challenges faced by practitioners in a certain field. It should be acknowledged that it can not cover all the questions posed by practitioners in the field of PTMs. However, as the most popular and representative QA platform in the programming community, leveraging SO’s extensive knowledge allows us to gain valuable insights into the PTM domain and contribute to existing research. Our findings can serve as a foundation for future work.

Identification of PTM-related questions. To include PTM-related questions, we created a set of PTM tags. However, this tag set may introduce bias as it may not cover all PTM-related questions. To address this limitation, we used a two-step approach. Firstly, we followed the methodology of previous work [13] to create an initial set of representative tags. Secondly, we used a snowball sampling technique, leveraging SO’s relevant tag recommendation mechanism, to expand our tag set. This iterative process helped reduce the risk of overlooking important PTM-related questions and ensure a more comprehensive coverage of the topic. However, it is worth noting that this approach does not guarantee that all identified questions are PTM-related. As shown in RQ3, around 5.81% of the questions were unrelated to PTMs. This potential threat may impact the results of RQ1&2.

Sampling of questions for manual analysis. To select questions for manual analysis in RQ3, we combine random sampling and high-score sampling strategies. While high-score sampling aimed to include questions of high quality and significance, the threshold ($score \geq 9$) for determining high scores is subjective to some extent. This trade-off between the representativeness of challenges and the feasibility of manual analysis is inevitable. It is important to note that the inclusion of high-score questions may slightly deviate from the real-world distribution of challenges.

Subjectivity of inspection. The results of RQ3 may be threatened by the manual analysis process. As with any qualitative research, there is a degree of subjectivity involved in interpreting and categorizing the data [88]. To mitigate this threat, we employed a rigorous and systematic approach. We established a codebook and clear guidelines for categorization, and the second and third authors were involved to ensure inter-rater reliability. We conducted regular discussions to address any disagreements or uncertainties.

X. CONCLUSION

This paper presents a comprehensive and systematic analysis of the challenges associated with PTM-related questions on SO. Our findings reveal that PTM-related questions are gaining increasing popularity and are difficult to solve. To gain deeper insights into these challenges, we construct a hierarchical taxonomy consisting of 42 codes and three categories, which provides a structured framework for understanding and addressing the specific challenges encountered in PTMs. Based on our findings, we discuss the implications of our study

for various stakeholders, including researchers, practitioners, vendors, and educators in the field of PTMs. To facilitate replication and further research, we make the data available online at <https://figshare.com/s/ef756cf1384e4542ddcb>.

XI. ACKNOWLEDGEMENTS

This work is supported by the National Natural Science Foundation of China Grants 62202022, 62302021, 62141209, and 62332001, and Self-determined Research Funds of State Key Laboratory of Complex & Critical Software Environment SKLSDE-2023ZX-15.

REFERENCES

- [1] C. Deguang, M. Jinlin, M. Ziping, and Z. Jie, "Review of pre-training techniques for natural language processing," *Journal of Frontiers of Computer Science & Technology*, vol. 15, no. 8, p. 1359, 2021.
- [2] B. Min, H. Ross, E. Sulem, A. P. B. Veysch, T. H. Nguyen, O. Sainz, E. Agirre, I. Heintz, and D. Roth, "Recent advances in natural language processing via large pre-trained language models: A survey," *ACM Computing Surveys*, 2021.
- [3] A. Karmakar and R. Robbes, "What do pre-trained code models know about code?" in *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 2021, pp. 1332–1336.
- [4] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [5] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever *et al.*, "Improving language understanding by generative pre-training," 2018.
- [6] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *The Journal of Machine Learning Research*, vol. 21, no. 1, pp. 5485–5551, 2020.
- [7] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. d. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman *et al.*, "Evaluating large language models trained on code," *arXiv preprint arXiv:2107.03374*, 2021.
- [8] E. Kalliamvakou. (2022, Sep.) Research: quantifying github copilot's impact on developer productivity and happiness. [Online]. Available: <https://github.blog/2022-09-07-research-quantifying-github-copilots-impact-on-developer-productivity-and-happiness/>
- [9] X. Han, Z. Zhang, N. Ding, Y. Gu, X. Liu, Y. Huo, J. Qiu, Y. Yao, A. Zhang, L. Zhang *et al.*, "Pre-trained models: Past, present and future," *AI Open*, vol. 2, pp. 225–250, 2021.
- [10] C. Zhou, Q. Li, C. Li, J. Yu, Y. Liu, G. Wang, K. Zhang, C. Ji, Q. Yan, L. He *et al.*, "A comprehensive survey on pretrained foundation models: A history from bert to chatgpt," *arXiv preprint arXiv:2302.09419*, 2023.
- [11] C. Treude and M. P. Robillard, "Augmenting api documentation with insights from stack overflow," in *Proceedings of the 38th International Conference on Software Engineering*, 2016, pp. 392–403.
- [12] M. U. Haque, L. H. Iwaya, and M. A. Babar, "Challenges in docker development: A large-scale study using stack overflow," in *Proceedings of the 14th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*, 2020, pp. 1–11.
- [13] Z. Chen, Y. Cao, Y. Liu, H. Wang, T. Xie, and X. Liu, "A comprehensive study on challenges in deploying deep learning based software," in *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2020, pp. 750–762.
- [14] C. Wang, Z. Chen, and M. Zhou, "Automl from software engineering perspective: Landscapes and challenges," in *Proceedings of the 20th International Conference on Mining Software Repositories*. MSR, 2023.
- [15] D. S. Cruzes and T. Dyba, "Recommended steps for thematic synthesis in software engineering," in *2011 international symposium on empirical software engineering and measurement*. IEEE, 2011, pp. 275–284.
- [16] H. Wang, J. Li, H. Wu, E. Hovy, and Y. Sun, "Pre-trained language models and their applications," *Engineering*, 2022.
- [17] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, "Language models are few-shot learners," *CoRR*, vol. abs/2005.14165, 2020. [Online]. Available: <https://arxiv.org/abs/2005.14165>
- [18] A. Oussidi and A. Elhassouny, "Deep generative models: Survey," in *2018 International conference on intelligent systems and computer vision (ISCV)*. IEEE, 2018, pp. 1–8.
- [19] X. Liu, F. Zhang, Z. Hou, L. Mian, Z. Wang, J. Zhang, and J. Tang, "Self-supervised learning: Generative or contrastive," *IEEE transactions on knowledge and data engineering*, vol. 35, no. 1, pp. 857–876, 2021.
- [20] J. Yang, G. Xiao, Y. Shen, W. Jiang, X. Hu, Y. Zhang, and J. Peng, "A survey of knowledge enhanced pre-trained models," *arXiv preprint arXiv:2110.00269*, 2021.
- [21] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz *et al.*, "Huggingface's transformers: State-of-the-art natural language processing," *arXiv preprint arXiv:1910.03771*, 2019.
- [22] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, "fairseq: A fast, extensible toolkit for sequence modeling," *arXiv preprint arXiv:1904.01038*, 2019.
- [23] H. Zhang, S. Wang, T.-H. Chen, Y. Zou, and A. E. Hassan, "An empirical study of obsolete answers on stack overflow," *IEEE Transactions on Software Engineering*, vol. 47, no. 4, pp. 850–862, 2019.
- [24] Wikipedia, "Stack Overflow — Wikipedia, the free encyclopedia," <http://en.wikipedia.org/w/index.php?title=Stack%20Overflow&oldid=1163163467>, 2023, [Online; accessed 10-July-2023].
- [25] P. Biernacki and D. Waldorf, "Snowball sampling: Problems and techniques of chain referral sampling," *Sociological methods & research*, vol. 10, no. 2, pp. 141–163, 1981.
- [26] S. Ahmed and M. Bagherzadeh, "What do concurrency developers ask about? a large-scale study using stack overflow," in *Proceedings of the 12th ACM/IEEE international symposium on empirical software engineering and measurement*, 2018, pp. 1–10.
- [27] C. Rosen and E. Shihab, "What are mobile developers asking about? a large scale study using stack overflow," *Empirical Software Engineering*, vol. 21, pp. 1192–1223, 2016.
- [28] X.-L. Yang, D. Lo, X. Xia, Z.-Y. Wan, and J.-L. Sun, "What security questions do developers ask? a large-scale study of stack overflow posts," *Journal of Computer Science and Technology*, vol. 31, pp. 910–924, 2016.
- [29] B. E. Storer and C. Kim, "Exact properties of some exact test statistics for comparing two binomial proportions," *Journal of the American Statistical Association*, vol. 85, no. 409, pp. 146–155, 1990.
- [30] M. Bagherzadeh and R. Khatchadourian, "Going big: a large-scale study on what big data developers ask," in *Proceedings of the 2019 27th ACM joint meeting on european software engineering conference and symposium on the foundations of software engineering*, 2019, pp. 432–442.
- [31] A. Baltadzhieva and G. Chrupala, "Predicting the quality of questions on stackoverflow," in *Proceedings of the international conference recent advances in natural language processing*, 2015, pp. 32–40.
- [32] B. Ci and R.-O. Rule, "Confidence intervals," *Lancet*, vol. 1, no. 8531, pp. 494–7, 1987.
- [33] P. I. Fusch Ph D and L. R. Ness, "Are we there yet? data saturation in qualitative research," 2015.
- [34] H. Zhao, Y. Li, F. Liu, X. Xie, and L. Chen, "State and tendency: an empirical study of deep learning question&answer topics on stack overflow," *Science China Information Sciences*, vol. 64, pp. 1–23, 2021.
- [35] Merle, "Disable layers in gpt-2 model," Apr. 2023. [Online]. Available: <https://stackoverflow.com/questions/76102276>
- [36] O. Ivanytskyi, "Cased vs uncased bert models in spacy and train data," May 2020. [Online]. Available: <https://stackoverflow.com/questions/61902426>
- [37] L. Elser, "Build an ensemble learning model for image multi-classification from pre-training," Dec. 2020. [Online]. Available: <https://stackoverflow.com/questions/65465801>
- [38] S. Silva, "Is it possible to view the code of a torch pretrained network," Apr. 2018. [Online]. Available: <https://stackoverflow.com/questions/50073402>
- [39] kloop, "How do i use bytelevelbptokenizer with utf-8?" Jun. 2022. [Online]. Available: <https://stackoverflow.com/questions/72761989>
- [40] R. Naushad, "Valueerror: Textcodeinput must be union[inputsequence, tuple[inputsequence, inputsequence]] - tokenizing bert / distilbert error," Aug. 2020. [Online]. Available: <https://stackoverflow.com/questions/63517293>

- [41] Y. Li and T. Yang, "Word embedding for understanding natural language: a survey," *Guide to big data applications*, pp. 83–104, 2018.
- [42] somethingstrang, "How to cluster similar sentences using bert," Apr. 2019. [Online]. Available: <https://stackoverflow.com/questions/55619176>
- [43] G. Cooper, "Using bert embeddings for seq2seq model building," Mar. 2020. [Online]. Available: <https://stackoverflow.com/questions/60927739>
- [44] O. Ivanytskyi, "Spacy's bert model doesn't learn," May 2020. [Online]. Available: <https://stackoverflow.com/questions/61943409>
- [45] Skalonga, "Data preprocessing for bert (base-german)," Jun. 2020. [Online]. Available: <https://stackoverflow.com/questions/62578609>
- [46] anegru, "Why is a throw-away column required in bert format?" Apr. 2019. [Online]. Available: <https://stackoverflow.com/questions/55910635>
- [47] Tia, "Addition of data to pre-trained model," Oct. 2018. [Online]. Available: <https://stackoverflow.com/questions/52777644>
- [48] user2182857, "What is the meaning of the second output of huggingface's bert?" Feb. 2020. [Online]. Available: <https://stackoverflow.com/questions/60243099>
- [49] Kaleem, "How to get all documents per topic in bertopic modeling," Oct. 2021. [Online]. Available: <https://stackoverflow.com/questions/69740911>
- [50] M. Manthiramoorthi, "How to train bart for text summarization using custom dataset?" Apr. 2020. [Online]. Available: <https://stackoverflow.com/questions/61390271>
- [51] Nilou, "How to use fine-tuned bert model in other tasks?" Feb. 2021. [Online]. Available: <https://stackoverflow.com/questions/66204936>
- [52] beginner, "Albert not converging - huggingface," Jun. 2020. [Online]. Available: <https://stackoverflow.com/questions/62490113>
- [53] k.avinash, "Hugging face transformer multiprocessing training," Jun. 2020. [Online]. Available: <https://stackoverflow.com/questions/62303645>
- [54] user901010, "Tensorflow hub directory not found when caching modules," Jun. 2018. [Online]. Available: <https://stackoverflow.com/questions/50808182>
- [55] kalsi, "Low ram usage & high gpu usage, hf datasets not helping in fine-tuning lm," Jan. 2023. [Online]. Available: <https://stackoverflow.com/questions/75065273>
- [56] F. Thayattchira, "Tensor flow model saving and calculating average of models," Jun. 2020. [Online]. Available: <https://stackoverflow.com/questions/62439765/>
- [57] S. Lee, "Save only best weights with huggingface transformers," Jun. 2020. [Online]. Available: <https://stackoverflow.com/questions/62525680/>
- [58] user15283025, "How to test a model before fine-tuning in pytorch lightning?" Sep. 2021. [Online]. Available: <https://stackoverflow.com/questions/69249187>
- [59] DevPy, "Improve the model prediction time in huggingface transformer models without gpu," Nov. 2021. [Online]. Available: <https://stackoverflow.com/questions/70077032>
- [60] A. Memon, "Convert a bert model to tflite," Apr. 2020. [Online]. Available: <https://stackoverflow.com/questions/60967842>
- [61] corvusMidnight, "Openai chatgpt (gpt-3.5) api error: 'invalidrequestor: Unrecognized request argument supplied: messages'," Mar. 2023. [Online]. Available: <https://stackoverflow.com/questions/75617865>
- [62] DanielBell99, "TypeError: not a string — parameters in autotokenizer.from_pretrained()," Jan. 2022. [Online]. Available: <https://stackoverflow.com/questions/70709572>
- [63] Overstack, "How do you implement openai gpt-3 api client in php?" Sep. 2022. [Online]. Available: <https://stackoverflow.com/questions/73794764>
- [64] Ljyeon, "How to fix python pip install openai error: subprocess-exited-with-error," Jan. 2023. [Online]. Available: <https://stackoverflow.com/questions/75067851>
- [65] Nyxnyx, "Huggingface error during training: Attributeerror: 'str' object has no attribute 'size'," Jan. 2021. [Online]. Available: <https://stackoverflow.com/questions/66834205>
- [66] Laz, "Transformer: cannot import name 'automodelwithlmhead' from 'transformers'," Sep. 2020. [Online]. Available: <https://stackoverflow.com/questions/64112358>
- [67] Code_7, "Multiprocessing/multithreading for huggingface pipeline," Aug. 2021. [Online]. Available: <https://stackoverflow.com/questions/68860984>
- [68] Exploring, "How to customize the positional embedding?" Aug. 2022. [Online]. Available: <https://stackoverflow.com/questions/73243093>
- [69] B. Bob, "Openai 'we could not parse the json body of your request.," Dec. 2022. [Online]. Available: <https://stackoverflow.com/questions/74903974>
- [70] M. Mutha, "How do i train a encoder-decoder model for a translation task using hugging face transformers?" Jun. 2020. [Online]. Available: <https://stackoverflow.com/questions/62446827/>
- [71] M. Allamanis and C. Sutton, "Why, when, and what: analyzing stack overflow questions by topic, type, and code," in *2013 10th Working conference on mining software repositories (MSR)*. IEEE, 2013, pp. 53–56.
- [72] P. S. Kochhar, "Mining testing questions on stack overflow," in *Proceedings of the 5th International Workshop on Software Mining*, 2016, pp. 32–38.
- [73] F. Fischer, K. Bottinger, H. Xiao, C. Stransky, Y. Acar, M. Backes, and S. Fahl, "Stack overflow considered harmful? the impact of copy&paste on android application security," in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, p. 121–136.
- [74] Y. Acar, M. Backes, S. Fahl, D. Kim, M. L. Mazurek, and C. Stransky, "You get where you're looking for: The impact of information sources on code security," in *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2016, pp. 289–305.
- [75] Z. Chen, H. Yao, Y. Lou, Y. Cao, Y. Liu, H. Wang, and X. Liu, "An empirical study on deployment faults of deep learning based mobile applications," in *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. IEEE, 2021, pp. 674–685.
- [76] A. Fontão, B. Ábia, I. Wiese, B. Estácio, M. Quinta, R. P. d. Santos, and A. C. Dias-Neto, "Supporting governance of mobile application developers from mining and analyzing technical questions in stack overflow," *Journal of Software Engineering Research and Development*, vol. 6, pp. 1–34, 2018.
- [77] L. Ponzanelli, A. Mocci, A. Bacchelli, M. Lanza, and D. Fullerton, "Improving low quality stack overflow post detection," in *2014 IEEE international conference on software maintenance and evolution*. IEEE, 2014, pp. 541–544.
- [78] G. Blanco, R. Pérez-López, F. Fdez-Riverola, and A. M. G. Lourenço, "Understanding the social evolution of the java community in stack overflow: A 10-year study of developer interactions," *Future Generation Computer Systems*, vol. 105, pp. 446–454, 2020.
- [79] I. Tenney, D. Das, and E. Pavlick, "Bert rediscovers the classical nlp pipeline," *arXiv preprint arXiv:1905.05950*, 2019.
- [80] M. Ballout, U. Krumnack, G. Heidemann, and K.-U. Kühnberger, "Opening the black box: Analyzing attention weights and hidden states in pre-trained language models for non-language tasks," *arXiv preprint arXiv:2306.12198*, 2023.
- [81] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut, "Albert: A lite bert for self-supervised learning of language representations," *arXiv preprint arXiv:1909.11942*, 2019.
- [82] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: a robustly optimized bert pretraining approach (2019)," *arXiv preprint arXiv:1907.11692*, vol. 364, 1907.
- [83] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter," *arXiv preprint arXiv:1910.01108*, 2019.
- [84] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman, "Glue: A multi-task benchmark and analysis platform for natural language understanding," *arXiv preprint arXiv:1804.07461*, 2018.
- [85] N. S. Keskar, B. McCann, L. R. Varshney, C. Xiong, and R. Socher, "Ctrl: A conditional transformer language model for controllable generation," *arXiv preprint arXiv:1909.05858*, 2019.
- [86] W. Yoon, J. Lee, D. Kim, M. Jeong, and J. Kang, "Pre-trained language model for biomedical question answering," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2019, pp. 727–740.
- [87] A. Abdellatif, D. Costa, K. Badran, R. Abdalkareem, and E. Shihab, "Challenges in chatbot development: A study of stack overflow posts," in *Proceedings of the 17th international conference on mining software repositories*, 2020, pp. 174–185.
- [88] C. Ratner *et al.*, "Subjectivity and objectivity in qualitative methodology," in *Forum Qualitative Sozialforschung/Forum: Qualitative Social Research*, vol. 3, no. 3, 2002.