

User Modeling in the Era of Large Language Models: Current Research and Future Directions

Zhaoxuan Tan, Meng Jiang

Department of Computer Science and Engineering, University of Notre Dame
{ztan3, mjiang2}@nd.edu

Abstract

User modeling (UM) aims to discover patterns or learn representations from user data about the characteristics of a specific user, such as profile, preference, and personality. The user models enable personalization and suspiciousness detection in many online applications such as recommendation, education, and healthcare. Two common types of user data are text and graph, as the data usually contain a large amount of user-generated content (UGC) and online interactions. The research of text and graph mining is developing rapidly, contributing many notable solutions in the past two decades. Recently, large language models (LLMs) have shown superior performance on generating, understanding, and even reasoning over text data. The approaches of user modeling have been equipped with LLMs and soon become outstanding. This article summarizes existing research about how and why LLMs are great tools of modeling and understanding UGC. Then it reviews a few categories of large language models for user modeling (LLM-UM) approaches that integrate the LLMs with text and graph-based methods in different ways. Then it introduces specific LLM-UM techniques for a variety of UM applications. Finally, it presents remaining challenges and future directions in the LLM-UM research. We maintain the reading list at: <https://github.com/TamSiuhin/LLM-UM-Reading>.

1 Introduction

User Modeling (UM) aims to extract valuable insights and patterns from user behaviors, enabling customization and adaptation of systems to meet specific users' needs [124]. UM techniques have facilitated a better understanding of user behaviors, customized intelligent assistance, and greatly improved user experience. For example, when people are looking for dinner options and searching online, the UM techniques infer their characteristics based on interaction history, predict current food interests, and give personalized recommendations. UM has a substantial impact on user data analysis and many applications such as e-commerce [191, 194, 280], entertainment [11, 33, 155], and social networks [1, 2, 212]. UM is a highly active and influential research field.

User modeling is mainly about mining and learning user data, including user-generated content (UGC) and user's interactions with other users and items. User-generated content encompasses a wide range of text data, such as tweets, reviews, blogs, and academic papers. The rich texts can be analyzed by natural language processing (NLP) techniques. User interactions, on the other hand, involve various actions such as following, sharing, rating, comments, and retweets. These interactions may form a heterogeneous temporal text-attributed graph [199] for having temporal and textual information and having different types of nodes and relations. That can be analyzed using graph mining and learning techniques. As a result, user modeling has branched out into text-based and graph-based approaches, focusing on extracting insights from text and graph data, respectively.

How has the research of text-based UM developed? Researchers have used multiple types of representations of texts, such as words, topics, and embeddings. Bag-of-Words (BoW) models create distributional text representations with word frequencies using a discrete vocabulary [67]. To address the sparsity of BoW representations,

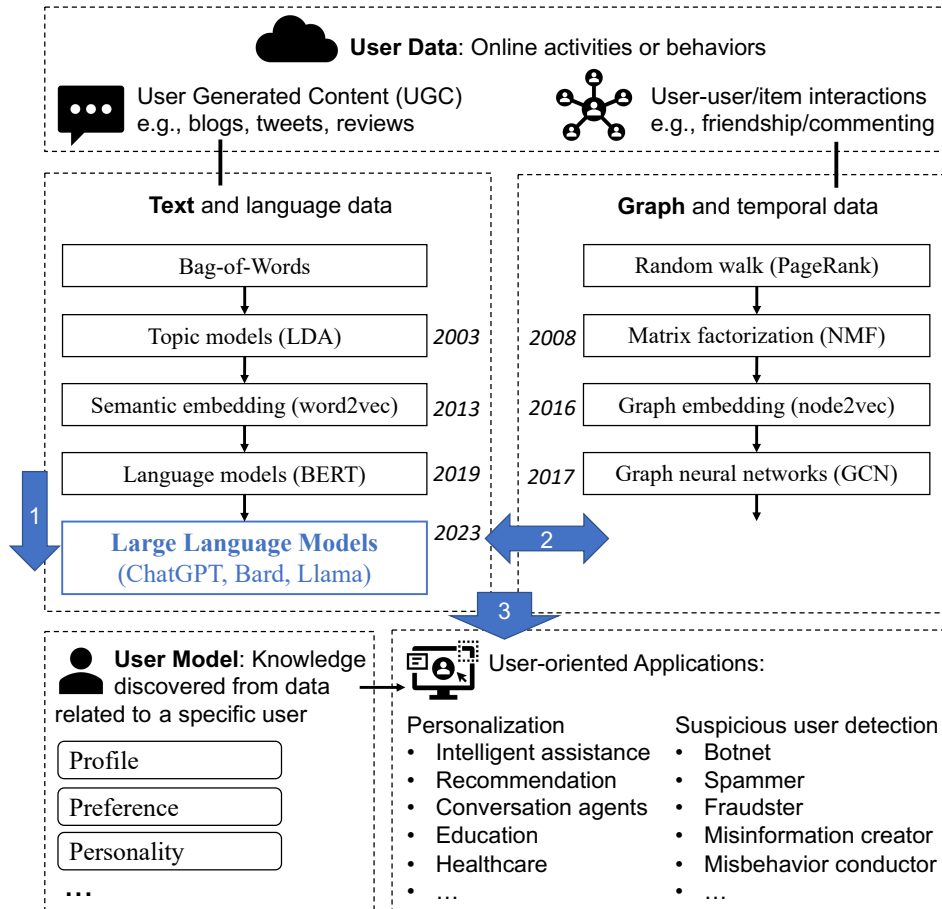


Figure 10: User modeling aims to discover knowledge and patterns from user data to identify profile, preference, and personality. The three blue arrows in the figure correspond to our three major contributions: (1) summarize how and why LLMs are great tools for modeling and understanding UGC, (2) review approaches that integrate LLMs with text- and graph-based UM methods, and (3) introduce LLM-UM techniques for various applications.

topic modeling techniques statistically discover latent topics in a collection of documents, e.g., latent Dirichlet allocation (LDA) [15]. But they are not able to capture semantic meanings, i.e., word semantic similarity. Word2Vec employs nonlinear neural layers to develop Continuous Bag-of-Words (CBOW) and continuous skip-gram models [148]. It extracts semantic embeddings from many kinds of UGC text data, such as blogs, reviews, and tweets. However, the neural layers are too shallow to capture deep sequential patterns among numerous word tokens. With the breakthrough in Transformer architectures [219], pretrained language models (PLMs) significantly changed the landscape of UGC understanding with the pretrain-fine-tune paradigm. The new paradigm trains the models on a large unlabeled corpus using self-supervised learning and uses hundreds or thousands of examples to fine-tune the models for downstream task adaptation [102]. Recently, large language models (LLMs) have revolutionized this area with emergent abilities, including unprecedented reasoning [241, 256], generalization [181, 239], and knowledge comprehension [163, 205]. LLMs are operated under the pretraining paradigm on extremely large corpora to update billions of parameters. A lot of research has shown that LLMs understand UGC in a zero-shot manner, i.e., without a collection of examples for fine-tuning. LLMs have surpassed human performance in summarization [176], outperformed most human in several exams [159], and shown strong reasoning abilities with prompt engineering, including Chain-of-Thought [241], Least-to-Most [290], and Tree-of-Thought [256]. LLMs open a new era for UM research to re-think about UGC mining.

How has the research of graph-based UM developed? User interactions with online content and users are naturally defined as edges that connect nodes of users or things. The user data can be defined as a graph. Heterogeneous graphs contain multiple types of nodes (e.g., users, items, places) and relations. Temporal/weighted graphs have timestamps/weights labeled on the interactions. Attributed graphs allow nodes to have a set of attribute-value pairs (e.g., age of a user, color of a product). In text-rich graphs, the nodes have long-form text attributes. Random walk with restarts provides a closeness score between two nodes in a weighted graph, and it has been successfully used in numerous settings (e.g., personalized PageRank [162]). Matrix factorization (MF) decomposes the user-item interaction matrix into the product of two matrices or known as latent features of users and items [91, 94, 107]. Regarding collaborative filtering, MF performs better with explicit feedback ratings, while RWR exploits the global popularity of items. It is actually a basic embedding model [272]. With deep learning, Node2Vec extracts sequences from the graph with random walks and uses Word2Vec to learn node embeddings [65]. However, encoding a graph into sequences would cause information loss. Graph neural networks (GNNs) employ the message-passing mechanism for deep representation learning on graphs. Specifically, the family of Graph Convolutional Network (GCN) [105] has greatly improved the performance in recommendation [49, 69], user profiling [24, 250], user behavior prediction [224, 263], and suspicious user detection [54, 55].

Why are LLMs revolutionizing text- and graph-based UM research? User modeling involves a series of machine learning tasks on text and graph data, such as text classification, node classification, link prediction, and time series modeling. Putting into context, the tasks can be sentiment analysis, natural language inference (NLI), user and product categorization, social relationship prediction, and temporal behavior prediction. Traditionally, the solution must be a specific model for a specific type of data and be trained on a specific set of annotations. For example, due to schema differences, two text classifiers had to be trained for the sentiment analysis and NLI tasks separately. Also, two networks or at least two modules in a graph neural network (GNN) were trained to predict if a user makes a new friend and purchases an item, respectively. Moreover, the textual information of user and/or product profiles is quite limited for learning and prediction due to the long-tail distribution.

LLMs have changed the paradigm of solution development. First, if designed properly, the prompts are able to treat most of the text-to-label tasks in LLMs as a unified text generation task; annotation data become not desperately needed; and the performance can even be comparable to or better than the traditional models. This is because LLMs were pre-trained on extremely large corpora and fine-tuned to follow the instructions in the prompts. Second, the prompts can be designed for learning tasks on graph data. For example, one can ask LLMs “if a user bought an Apple watch yesterday, will the user consider buying a pair of running shoes?” The “analysis” by LLMs can provide additional information to existing user-item link predictors. Third, all text information can be automatically expanded by LLMs. The relevant parametric knowledge augments the input of machine learning models and reduces the task difficulty.

LLMs have showcased robust capabilities in characterizing users’ personalities [184], discerning users’ stances [271], pinpointing user preferences [52], and beyond. Also, they have demonstrated marked proficiency in node classification [259], node attribute prediction [70], and graph reasoning [226]. Preliminary research focuses on leveraging LLMs for user modeling (LLM-UM) to integrate text-based and graph-based methods. For user profiling, GENRE [135] leverages ChatGPT as a user profiler by feeding users’ behavior history and prompting the model to infer the users’ preferred topics and regions. These LLM-generated profiles serve as important features for click-through rate recommendation models and resolve the anonymity problem in collecting user profiles. For recommendation, Kang et al. [100] use LLMs to predict user ratings based on their behavior history and find that LLMs typically demand less data while maintaining world knowledge about humans and items. For personalization, LaMP [189] proposes a benchmark incorporating personalized text generation and classification tasks as well as a retrieval-augmented approach. LLMs can be personalization tools for their understanding of user data. On suspiciousness detection, Chiu et al. [29] employ GPT-3 to detect hate speech, discovering that LLMs are able to identify abusive language with limited labels.

Difference from existing surveys. Given the growing interest and expanding body of work in user modeling (UM) with LLMs, this is an ideal opportunity to provide a comprehensive review that accomplishes several goals. In this survey, we analyze the advantages of LLMs in boosting existing user modeling techniques, introduce a taxonomy that categorizes LLM-UM techniques from the perspective of methodology, provide an in-depth review of specific techniques for a wide range of real-world applications, and finally outline the challenges and potential future directions in the field. We aim to provide a handbook for researchers and practitioners in the relevant fields, so they are able to confidently use LLMs to design and develop effective UM approaches.

It is worthwhile to discuss the uniqueness of our effort. Farid et al. [53] conducted a survey on user profiling approaches as a part of user modeling before LLMs came out. Li and Zhao [124] reviewed representation learning methods for user modeling, which was a prevailing paradigm before the advent of LLMs. He et al. [71] focused on user behavior modeling within recommender systems, narrowing down the scope specifically to item recommendations in the pre-LLM era. In the context of the LLM era, Fan et al. [50], Wu et al. [243] investigated the applications of LLMs in recommender systems. Additionally, Lin et al. [130] examined approaches that incorporated LLMs to enhance recommender systems, while Chen et al. [22] summarized recent work, challenges, and future directions in LLM-based personalization. However, these surveys discussed specific application goals of user modeling, either recommendation or personalization, which represents only a small portion of user modeling. To the best of our knowledge, there is no survey that summarizes the LLM-UM research work. Hence, our survey aims to fill this gap by providing a comprehensive summary and inspiring future research directions.

The remainder of this survey is organized as follows (Figure 11). Section 2 gives the background of user modeling techniques and large language models and gives the motivation of why LLMs are good tools for next-generation user modeling. Section 4 introduces two taxonomies of LLM-UM based on their approaches and applications. Section 5 summarizes the approaches to LLM-UM and how LLMs can integrate text and graph-based methods in existing works, including leveraging LLMs as enhancers, predictors, and controllers. Section 6 elaborates on the LLM-UM applications, including personalization and suspiciousness detection. Finally, Section 7 delves into current challenges and future directions pertaining to the LLM-UM topic.

2 Background

2.1 User Modeling

User modeling (UM) involves insight extraction or prediction from user data, such as profiles, personality traits, behavior patterns, and preferences. The insights can be utilized to customize and optimize user-oriented systems or services, enabling them to adapt effectively to the unique requirements of individual users [124]. From a user data standpoint, there are two primary types: user-generated content (UGC) and user-user/item interaction. These data modalities encompass textual content and graph-based interactions. User modeling techniques can be broadly classified into two categories: text-based methods and graph-based methods, each with a distinct emphasis on UGC and user interaction graphs, respectively.

Text-based UM Methods. The text-based UM methods focus on mining user-generated content, understanding user profiles, personalities, and subsequently inferring user preference and providing personalized recommendations and assistance. Text-based UM methods are highly relevant to natural language processing and have experienced several breakthroughs in the last two decades.

Initially, people widely use Bag-of-Words (BoW) [67] that maintains a discrete word vocabulary and represents text as an unordered collection of words. However, BoW disregards grammar, word order, and semantics. Later, topic models are proposed to identify groups of similar words in an unsupervised statistical machine learning manner [5]. The most representative topic model is Latent Dirichlet Analysis (LDA) [15], which considers documents as a mixture of topics and topics as a mixture of words. Latent semantic analysis (LSA) [47] is another popular topic model that is based on the principle that words close in meaning tend to be used together in context.

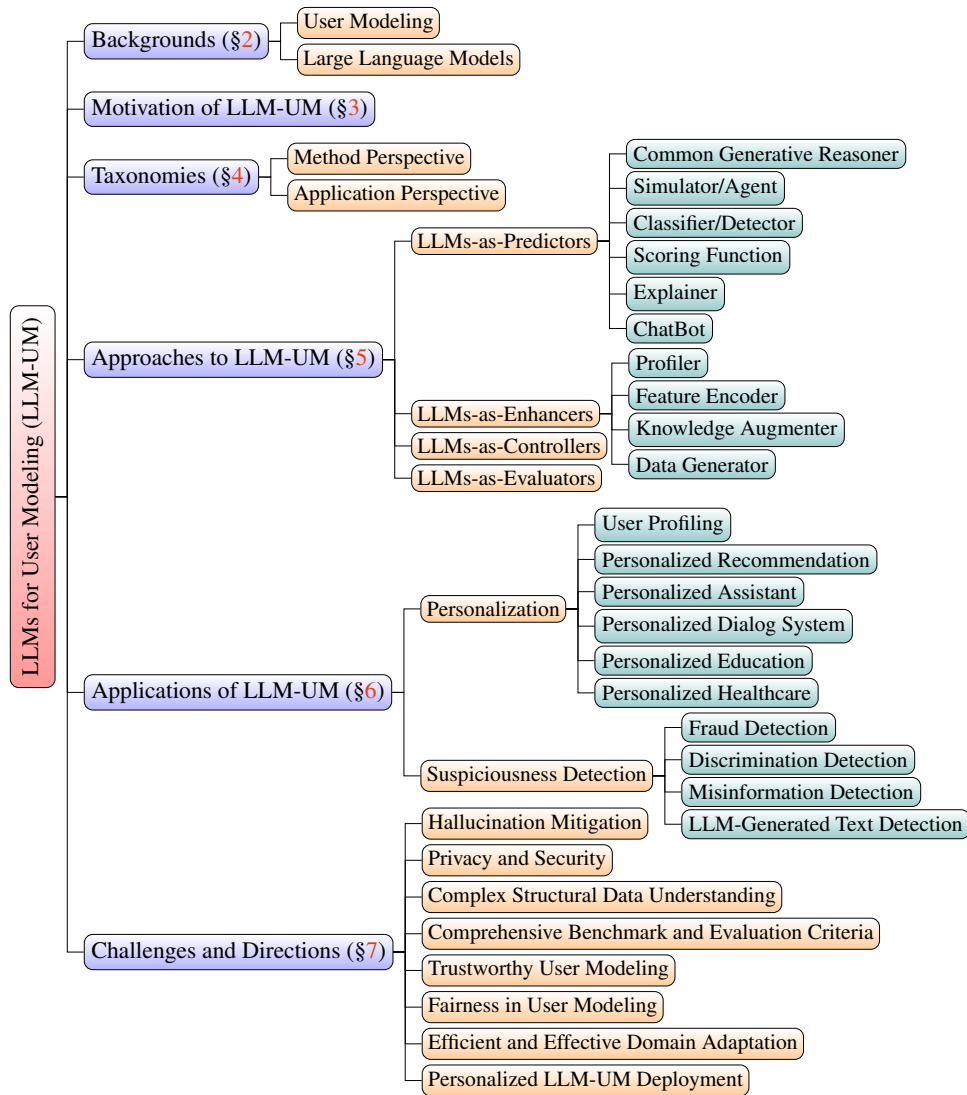


Figure 11: Structure of this survey.

Soon after, Mikolov et al. [148] introduced neural network techniques for UGC mining and understanding, including the continuous Bag-of-Words (CBOW) that predicts current words based on the context and continuous skip-gram model that predicts surrounding words given the current words. The emergence of Word2Vec brings new ideas to user modeling; for example, Hu et al. [75] use Word2Vec to encode user search history and predict users' age, gender, and education. However, the simplicity of Word2Vec leads to a limited semantic understanding of UGC. The emergence of Transformer architecture [219] and pretrained language models (PLMs) [102, 179] dramatically change the landscape of natural language processing, resulting in significant progress in text-based user modeling. For instance, BERT4Rec [204] utilizes a bidirectional self-attention mechanism to learn user behavior sequence representations for recommendation. Recently, large language models (LLMs) and their unprecedented emergent ability lead to a revolution in text-based user modeling. Specifically, LLMs are trained on large amounts of textual data with billions of learnable parameters to understand the patterns, semantics, and structure of natural language. Some preliminary explorations have shown LLMs are instrumental in enhancing functionality and adapting to users' specific needs in user modeling systems. LLMs have the capability to serve as a user profiler to generate user characteristics and personality, as in representative work

such as GENRE [135], PALR [26], and MBTI-based assessment Rao et al. [184]. Furthermore, LLMs are proven to be powerful recommender systems, evidenced by LLMRec [134], Chat-REC [63], and LKPNR [187]. LLM personalization has also been a prominent area of research, delivering products such as LaMP [189], AuthorPred [117], and NetGPT [25]. Additionally, several studies have explored the utilization of LLMs for the detection of misinformation and misbehavior [110, 152].

Graph-based UM Methods. The graph-based user modeling methods focus on learning from graph structures about user-user/item interactions. Moreover, the timestamp and intrinsic heterogeneity of interaction networks enrich the graph information, together forming a temporal heterogeneous graph. To mine the graph and temporal data for user modeling, PAGE [162] first propose PageRank, which measures the importance of a node based on times visited by the random walk. Collaborative filtering [193] is then proposed, assuming that users with similar behaviors would rate and act on items in a similar manner. Matrix factorization, as a collaborative filtering technique, has dominated the user modeling field for quite a few years since Netflix competition for its scalability and flexibility [107, 113, 202]. In its basic form, matrix factorization features both items and users in latent space based on user-item interaction history. High correspondence between the item and the user leads to a recommendation. Later, deep learning models are introduced to learn high-quality user and item embeddings in latent spaces [45, 65, 172, 209]. Node2Vec [65] is the pioneering work that uses random walks on social networks to generate sentences from graph structures and then feeds them into a Word2Vec model [148] for graph embedding learning. Metapath2Vec [45] extends Node2Vec to heterogeneous graphs. To mitigate Node2Vec’s information loss by extracting graphs into sequences, Graph Neural Network (GNN) is applied to advance user modeling with their robust structural encoding capabilities [77, 105, 220]. The core of graph neural networks lies in the message-passing mechanism that propagates node representations and aggregates neighborhood representations. The most representative GNN is graph convolutional network (GCN), which aggregates the information of different neighbors equally. Then the graph attention network (GAT) uses an attention mechanism to learn the attention weights of neighborhoods for higher-quality node representations. Specifically in user modeling applications, GNN-based methods are widely adopted in advanced user modeling systems and have achieved state-of-the-art performance. For instance, Ying et al. [261] design an efficient random walk algorithm, merge it with a type of GNN named GraphSAGE, and deploy the system on Pinterest.

2.2 Large Language Model

Language models are probabilistic models of natural language that can generate the likelihood of word sequences so as to predict the probabilities of future tokens [119, 282]. Large language models (LLMs) refer to the deep neural language models with billions of learnable parameters that are pretrained on an extremely large textual corpus to understand the distribution and structure of natural language [282]. Thanks to the efficiency of the Transformer architecture [219], almost all large language models employ it as the backbone. There are three types of language model design: encoder-only (e.g., BERT [102]), decoder-only (e.g., GPT [180]), and encoder-decoder (e.g., T5 [182]). Encoder-only models, specifically for BERT, use bidirectional attention to process token sequences and are pre-trained on masked token prediction and next-sentence classification tasks. This process can extract semantic embeddings for general purposes and enable the models to quickly adapt to diverse downstream tasks after fine-tuning. Decoder-only models, such as GPT, conduct text-to-text tasks based on the transformer decoder architecture. They are trained on the next token prediction tasks from left to right generation. Encoder-decoder models, such as T5, are trained on text-to-text tasks. Their encoders extract contextual representations from the input sequence, and their decoders use cross attention to map latent representations back to the text output space. In the context of LLMs, most models follow the decoder-only architecture as it simplifies the model and makes efficient inferences [232].

Recently, researchers have found that scaling pretrained language models’ training data and parameter size often leads to a significant performance gain, a.k.a scaling law [101]. The large language models present emergent abilities [240], referring to the abilities that are not present in small models. Typically, there are three types of

well-studied emergent abilities: in-context learning (ICL), instruction following, and step-by-step reasoning. In-context learning assumes that the language model has been provided with natural language instructions and/or several task demonstrations. LLMs can generate the expected output for test instances by completing the word sequence of input text without requiring additional training or gradient update, which is first introduced by GPT-3 [17]. Recent ICL research focuses on reducing inductive bias [116, 201]. The instruction following ability means that LLMs are shown to perform well on unseen tasks that are also described in the form of instructions after fine-tuning with a mixture of multi-task datasets formatted via natural language descriptions, known as instruction tuning. Instruction tuning improves the generalization ability of LLMs. The LLMs are better aligned with human intentions [238]. Recent instruction tuning studies focus on how to align LLMs with tasks and user preferences effectively [161] and efficiently [289]. Step-by-step reasoning means that LLMs can solve complex tasks that require multi-step reasoning. Chain-of-Thought (CoT) [241] introduces intermediate steps of reasoning steps in prompt design. Least-to-most [290] breaks down reasoning steps into simpler problems. Self-consistency [234] prompting further enhances LLMs reasoning by ensembling diverse CoT reasoning paths. Tree-of-Thought (ToT) [256] and Graph-of-Thought (GoT) [12, 257] enable LLMs to explore the thought processes in tree and graph structure, respectively. Moreover, preliminary explorations show that LLMs can use external tools [195], be parametric knowledge bases [163], have theory-of-mind [108], act as agents [230, 247], have graph understanding abilities [226], and can serve as optimizers [251].

Apart from using LLMs with frozen parameters, another line of work focuses on efficiently fine-tuning parameters in LLMs, namely parameter efficient fine-tuning, which helps LLMs efficiently adapt to specific tasks, datasets, and domain-specific understanding. Prefix tuning [126] keeps the language model parameters frozen and optimizes a small continuous task-specific vector called the prefix. Prompt tuning [115] is a simpler variant of prefix tuning, where some vectors are prepended at the beginning of a sequence at the input layer. Llama-Adapter [277] appends a set of learnable prompts as the prefix to the input instruction tokens in the higher transformer layers of Llama [216]. LoRA [74] injects trainable rank decomposition matrices into each layer of the transformer architecture, greatly diminishing the number of trainable parameters for downstream tasks. QLoRA [42] updates parameters through a frozen 4-bit LLM into low-rank adapters.

Existing LLMs can be categorized into open-source and API-based models based on accessibility. Open-source models provide access to both model weights and the ability to run the models on local machines, while API-based models restrict users from directly accessing the model weights and only allow them to interact with the models through an API. The open-source LLMs include T5 [182], Flan-T5 [34], OPT [278], BLOOM [192], GLM [270], Llama [216], and Falcon.¹ By fine-tuning or instruction-tuning Llama, a family of LLMs emerged, such as Alpaca [213] and Vicuna [28]. For API-based models, OpenAI offers four major series of GPT-3 [17] interface, including `ada`, `babbage`, `curie`, and `davinci`, corresponding to GPT-3 (350M), GPT-3 (1B), GPT-3 (6.7B), and GPT-3 (175B). GPT-3.5 series includes `davinci` and `turbo:turbo` leverages Reinforcement Learning from Human Feedback (RLHF) [161] to create human-like conversations. GPT-4 [159] is a well-acknowledged state-of-the-art and achieves astonishing performance on a wide range of tasks. There are some other API-based models such as BARD [146], Claude,² PaLM [31], BloombergGPT [245], and LangChain.³

Despite the thriving development of the LLM research, some challenges remain unsolved. For example, LLMs suffer from hallucination issues, that being said, LLMs generate text that is fluent and natural but unfaithful to the source content or under-determined [99]. LLMs also have biases due to unfathomable pre-training datasets, including political discourse [57], hate speech [81], and discrimination. Moreover, the inference latency of LLMs remains high because of low parallelizability and large memory footprints [175]. The remaining challenges also include limited context length [99], outdated knowledge [258], misaligned behavior [188], brittle evaluation

¹<https://falconnllm.tii.ae/our-research.html>

²<https://www.anthropic.com/index/claude-2>

³<https://www.langchain.com/>

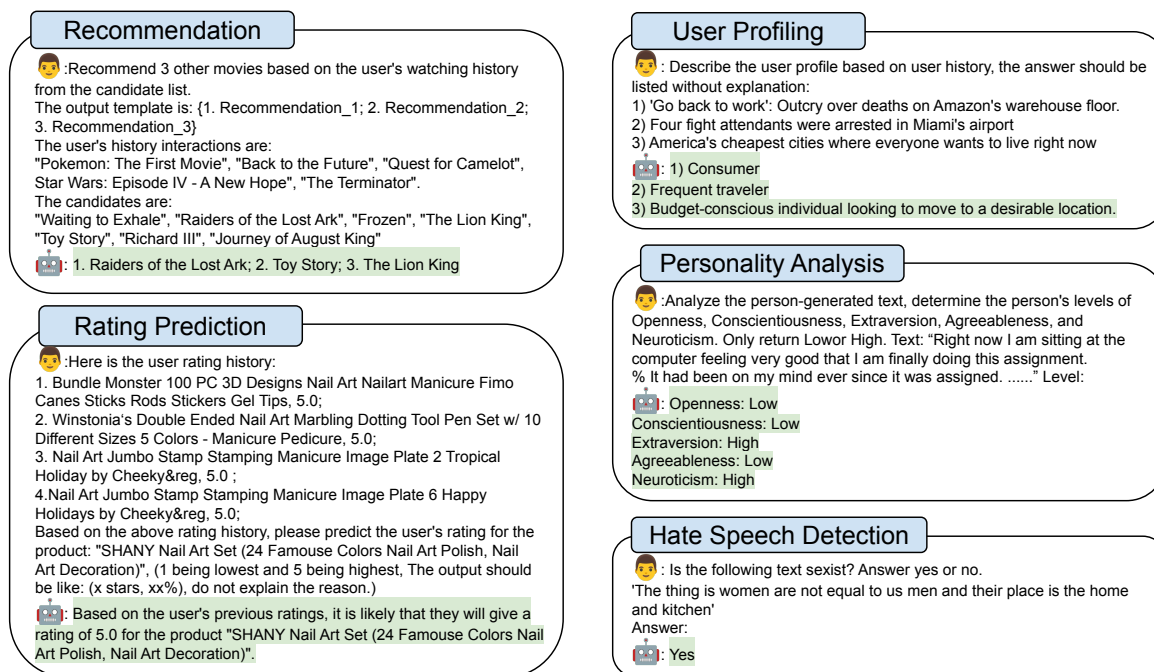


Figure 12: Some examples of LLMs for recommendation, rating prediction, user profiling, personality analysis, and hate speech detection. They serve as compelling examples that demonstrate the ability of LLMs to effectively model, comprehend, and reason based on user-generated content (UGC) and user interactions.

[283], and limited structure understanding ability [27].

3 Motivation: Why LLMs for User Modeling?

LLMs have demonstrated novel capabilities, exhibiting strong potential in modeling and understanding user-generated content (UGC). In this section, we present related studies and specific examples (Figure 12) to support the claim that LLMs are effective tools for UGC modeling and comprehension.

A growing body of research focuses on utilizing LLMs for recommendation purposes, aiming to predict users' item-based interests from their behavior history. For example, PALR [26] generates a user profile based on UGC history and constructs prompts that incorporate the history, profile, and item candidates to enable LLMs to provide recommendations. As illustrated in the real example depicted in Figure 12, LLMs successfully generate reasonable recommendations by considering the user's viewing history. In the domain of user profiling, the objective is to summarize users' characteristics, including personality, interests, and topics of interest, based on their generated content and history. Recent research demonstrates that LLMs excel in user profiling. Liu et al. [136] put behavior history into LLMs and extract users' interest topics and physical regions and thus augment recommender systems. The example in Figure 12 further confirms the effectiveness of LLMs in summarizing user characteristics, preferences, and intentions. In the context of rating prediction, which seeks to comprehend user preferences based on past UGC, Kang et al. [100] investigate the ability of LLMs to predict user ratings for candidate items. They find that zero-shot LLMs lag behind traditional recommendation models due to the absence of user interaction data. LLMs can employ reasoning based on users' previous ratings to predict ratings for candidate items. Recent studies highlight LLMs' capacity to understand user personality. Ji et al. [87] employ various prompts to probe LLMs' ability to recognize user personality based on UGC history. LLMs achieve impressive results in personality recognition under zero-shot conditions. For users with a harmless behavior

history, the UM is built upon general interests, content preferences, and interaction styles. However, the UM for users with suspicious behavior history, e.g., hate speech, would have to assess behavior patterns correlated to the suspicious activity, the risk of future incidents, and potentially flag these users for closer monitoring. From the application aspect, the presence or absence of hate speech in users' history has a significant impact on personalized recommendations. For example, the users with a hate speech history might be steered away from sensitive topics in recommendations. Therefore, suspiciousness detection, e.g., hate speech detection, is an essential application in user modeling. LLMs are good at detecting suspiciousness in UGC. Del Arco et al. [40] explore zero-shot prompting LLMs for hate speech detection. They find that zero-shot prompting can achieve performance comparable to and even surpass fine-tuned models. Figure 12 presents an example to further validate LLMs' effectiveness in detecting suspicious UGC.

Collectively, these studies and examples demonstrate LLMs' capabilities of modeling, understanding, and reasoning UGC and user behavior. They provide comprehensive evidence that LLMs can serve as valuable tools for user modeling, showcasing significant potential to improve user-oriented applications.

4 Taxonomies

In this section, we present taxonomies that classify LLM-UM techniques based on their approaches and applications. In Figure 11, we structure the survey according to these two taxonomies in Section 5 and 6, where the upper-level section is the parent category and the lower-level section is the child concept. For instance, the concept of "Approaches" to LLM-UM encompasses "LLMs-as-Predictors", "LLMs-as-Enhancers", "LLMs-as-Controllers", and "LLMs-as-Evaluators".

4.1 Approach Perspective

LLMs play diverse roles in LLM-UM systems. Based on their functionality, LLM-UM work can be categorized into four distinct approaches. (1) **LLMs-as-Predictors** involves leveraging LLMs for reasoning and generating answers directly. Depending on the role of LLMs, these LLM-UM methods can further be categorized as common generative reasoners for complex tasks, agents/simulators that model and predict human behavior, classifiers/detectors, scoring functions, explanation generation, and Chatbots for user modeling. (2) **LLMs-as-Enhancers** refers to using LLMs as augmentation modules to enhance the downstream user modeling system. LLMs can act as profilers to infer user preferences and characteristics, serve as feature encoders to generate latent UGC representations, augment discriminative user modeling systems with knowledge stored in LLMs, and generate high-quality data for small UM model training. (3) LLMs also possess the ability to control the pipeline of UM systems (**LLMs-as-Controllers**), automatically determining whether to execute certain operations. (4) LLMs can also serve as evaluators (**LLMs-as-Evaluators**) to score and analyze conversations and text generated under open-domain settings.

4.2 Application Perspective

Another taxonomy for categorizing LLM-UM is based on the downstream applications they address. Generally, LLM-UM systems aim to adapt to users' personal needs and detect suspiciousness in user data.

Personalization refers to tailoring experiences to individual preferences. Existing LLM-UM works can be categorized as follows: user profiling, which aims to tell the characteristics, personalities, stances, and sentiments towards certain targets; personalized recommendation gives the recommended items based on the user's behavior history and user profile; personalized assistants aim to adapt to users' specific needs and provide customized assistance; personalized dialog systems that interact iteratively with humans based on user data; and personalized applications in the education and healthcare domains. **Suspiciousness detection** is focused on identifying malicious users and UGC such as fraudsters, spammers, discriminations, and misinformation to preserve the

integrity of social discourse [95, 208]. Existing LLM-UM works in suspiciousness detection can be categorized as follows: fraud detection for identifying malicious users or information in social platforms; discrimination detection to combat the spread of hate speech, predatory, and sexist; misinformation detection to identify fake news and propaganda; and LLM-generated text detection to identify if text is AI-generated for the integrity of education, online discourse, etc.

5 Approaches to LLM-UM

Given the strong capabilities in generation [282], reasoning [241], knowledge comprehension [205], and a good understanding of UGC as elaborated in Section 3, LLMs can be used to supercharge UM systems. The LLM-UM approaches can be generally categorized into three categories based on the LLMs' role, where the first envisions LLMs as the sole predictor that generates prediction directly, the second employs LLMs as enhancers to probe more information for the UM system augmentation, the third empowers LLMs with the ability to control the UM methods pipeline, automating the UM process, and the last uses LLMs as evaluators, assessing the performance of the system. It's worth mentioning that the form of "user model" in LLM-UM remains consistent with the previous definition, encompassing the knowledge and patterns that are discovered with the help of user-generated content, and user-user/item interaction networks [71]. The distinction of LLM-UM from previous paradigms lies in the approaches, where LLM-UM are empowered or enhanced by LLMs to gain user-related knowledge. In the following subsections, we summarize each paradigm and present representative approaches.

5.1 LLMs-as-Predictors

In this section, we introduce the LLMs-as-Predictors paradigm presented in LLM-UM works, which means LLMs are leveraged to make predictions and generate answers for downstream applications directly. More specifically, approaches in leveraging LLMs as generative reasoners, simulators/agents, classifiers/detectors, scoring/ranking functions, explainers, and Chatbots.

5.1.1 Common Generative Reasoner

Thanks to the strong generalization of LLMs, numerous UM systems incorporate them as generative reasoners to directly execute downstream tasks with a method design that is free of shenanigans. Considering the significant body of work falling under the generative reasoner category, we further classify them into non-tuning and fine-tuning methods based on the parameter status of LLMs.

Non-tuned LLMs. Utilizing LLMs as non-fine-tuning generative reasoners involves directly employing pretrained LLMs without adjusting their parameters for specific tasks. This approach preserves their versatility and ability to generalize across diverse circumstances. Therefore, the focus of non-tune LLMs research is mainly on prompt template and pipeline design. Di Palma et al. [44] present a comprehensive experimental evaluation framework to examine LLMs' ability in recommendation rigorously. Liu et al. [133] investigate LLMs capability in recommendation by designing prompts to test ChatGPT on directly performing rating prediction, sequential recommendation, direct recommendation, explanation, generation, and review summarization. NIR [229] designs a three-stage prompt template that guides LLMs to understand user preference, select representative movie-watching history, and recommend a list of movies. Sanner et al. [190] explore the prompting LLMs with item-based and language-based preference and find LLMs are strong near cold-start recommendation systems for pure language-based preference. BookGPT [288] design prompt templates and the overall pipeline to perform book rating recommendation, user rating recommendation, and book summary recommendation. Hou et al. [73] design prompting template by including the sequential interaction history, candidate items, and ranking instruction and show that LLMs have promising zero-shot ranking abilities in recommendation tasks. PALR [26]

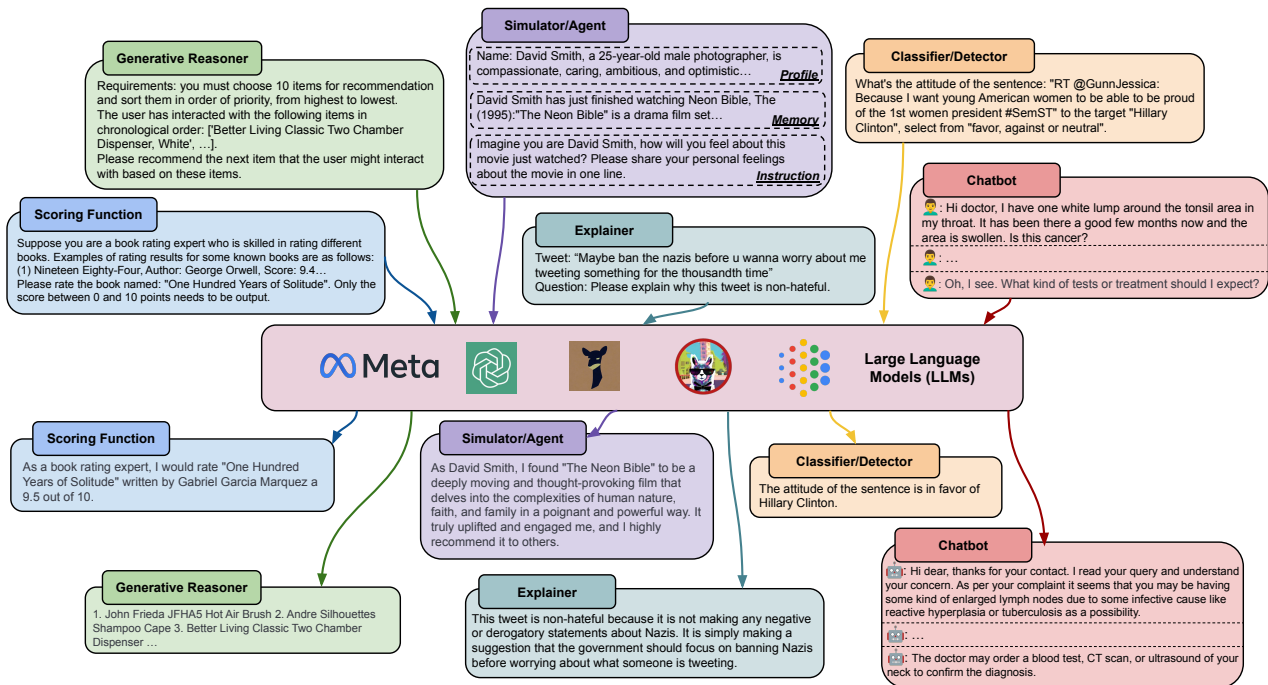


Figure 13: LLMs-as-Predictors, where LLMs are exclusively utilized to generate the predicted response.

designs prompt templates containing user interaction history sequences, LLM-generated user profiles, and candidates, then feeds them into a general-purpose LLM for recommendation. LaMP [189] offers an evaluation framework for personalized content generation and classification. Li et al. [117] design a multistage prompting strategy and multitasking to help frozen LLMs generate personalized content. The multistage contains retrieval, ranking, summarization, synthesis, and generation, and multitasking is to predict if the same author writes two documents. Li et al. [125] investigate ChatGPT’s capabilities in personalized news recommendation, news provider fairness, and fake news detection, where they find ChatGPT is sensitive to input phrasing. Rao et al. [184] explore LLMs capabilities to analyze human personalities based on the Myers-Briggs Type Indicator (MBTI) test. Ji et al. [87] investigate the text-based personality recognition ability of ChatGPT and propose a level-oriented prompting strategy to optimize zero-shot chain-of-thought performance on personality recognition. Ghanadian et al. [64] leverage ChatGPT to assess the suicide risks on social media with zero-shot and few-shot prompting. Fan and Jiang [51] investigate ChatGPT’s capabilities in discourse dialogue analysis, including topic segmentation, discourse relation recognition, and discourse parsing. Wu et al. [244] prompt LLMs to pairwise compare lawmakers and then scale the resulting graph using the Bradley-Terry model, finding that LLMs can be used to estimate the latent positions of politicians.

Fine-tuned LLMs. Fine-tuning LLMs can facilitate their adaptation to specific UM tasks and lead to better domain specialization. Given the massive number of parameters, research generally employs parameter-efficient fine-tuning (PEFT) techniques, such as LoRA, Llama-adapter, prompt-tuning, etc. InstructRec [276] considers LLMs as generative instruction following recommender and performs instruction tuning on Flan-T5-XL [34] model with a large amount of user-personalized instruction data. GIRL [287] proposes a Proximal Policy Optimization (PPO)-based reinforcement learning method to fine-tune LLMs, which aims to improve the LLM’s ability to assess the compatibility between a job and a user. TallRec [8] leverages rec tuning samples containing user history behavior, new items, and feedback for instruction tuning, and it uses LoRA to improve the efficiency as well. GLRec [242] constructs prompts based on meta paths extracted from the user behavior graph, then leverages weighted path embeddings, instruction tuning, and LoRA for LLMs tuning and generates recommended

items. Tie et al. [215] leverage fine-tuned LLMs to provide clinically useful, personalized impressions.

Combining non-tune and fine-tune paradigms can improve LLM performance. For example, LaMP [189] proposes a retrieval-augmented approach to retrieve personalized history to construct prompts for LLM generation under zero-shot and fine-tuning settings. Christakopoulou et al. [32] investigate the user’s interest journey using few-shot prompting, prompt-tuning, and fine-tuning for journey name extraction. ReLLa [131] proposes to retrieve semantic user behavior to augment LLMs under the zero-shot settings and design retrieval-enhanced instruction tuning for the few-shot recommendation.

5.1.2 Simulator/Agent

Using LLMs as autonomous agents has been a prosperous research direction recently, which expects LLMs to accomplish tasks through self-directed planning and actions [230]. In the user modeling domain, LLMs can serve as a user simulator to imitate user behavior, conduct UM applications with planning and actions, use external tools, etc. Some works leverage LLM as a user simulator to predict user behavior. RecLLM [61] plugs LLM into the conversational recommender system to generate synthetic conversations to simulate user behavior for tuning system modules. UGRO [78] uses LLM as an annotation-free user simulator to assess dialogue responses. Kong et al. [106] fine-tune LLMs on genuine human-machine conversations to get a user simulator to generate a high-quality human-centric synthetic conversation dataset. PersonaLLM [90] investigates whether the behavior of LLM-generated personas can reflect certain personality traits accurately and consistently.

A few research studies have envisioned LLMs as agents and enabled them to interact with and explore the environment to gain a better understanding of user modeling tasks. Generative Agent [166] leverages LLMs to simulate human behavior in a social context with a memory module to record agent history and memories, and retrieve them for planning. RecMind [236] presents an LLM-powered autonomous recommender agent capable of providing precise, personalized recommendations through careful planning by utilizing tools for obtaining external knowledge and leveraging individual data. InteRecAgent [80] employs LLMs as the brain and recommender models as tools. It comprises key components such as a memory bus, dynamic demonstration-augmented task planning, and reflection. RecAgent [231] constructs a user simulator that regards each user as an LLM-based autonomous agent and lets different agents freely communicate, behave, and evolve within the recommender system. AutoGPT⁴ demonstrates autonomous comprehension of specific objectives using natural language and carries out automated processes in a continuous loop, effectively accomplishing user-specific tasks. LLMs as agents can also empower LLMs with external tools and API for better user understanding. For example, Graph-Toolformer [273] teaches LLMs to use external graph-related API to augment LLMs’ ability to reason over structural data and, therefore perform sequential recommendation and user rating prediction.

5.1.3 Classifier/Detector

LLMs can be prompted to serve as classifiers and detectors to analyze UGC, e.g., detect stance, hate speech, and suspicious behavior [186, 206]. Zhang et al. [271] discuss the potential of LLMs in stance detection and explanation generation. Hu et al. [76] propose ladder-of-thought (LoT) that assimilates high-quality external knowledge in small LMs to augment stance detection in LLMs. Mu et al. [153] investigate LLMs potential in zero-shot text classification in computational social science, including bragging, vaccine, complaint, and hate speech detection. Ziems et al. [292] comprehensively investigate LLMs capabilities in classifying dialect features, emotion, humor, empathy, discourse acts, and more detection and classification tasks. Parikh et al. [165] explores LLMs capabilities in user intent classification under zero-shot and few-shot settings. Qin et al. [178] construct an interactive depression detection system based on LLMs with CoT and a selector of users’ tweets. Ferrara [58] discuss social bot detection in the era of LLMs, which points out that LLMs can be used to improve bot detection in low-resource language and domain. Qi et al. [177] design zero-shot, few-shot, and fine-tuning

⁴<https://news.agpt.co/>

methods to perform suicidal risk and cognitive distortion identification on Chinese social media. SentimentGPT [104] investigates the GPT’s sentiment analysis capabilities in prompt engineering and fine-tuning. ALEX [97] prompts an LLM with the predicted results from BERT models for health data classification. Finch et al. [59] leverage GPT-3.5 to perform dialogue behavior detection for nine categories. Bhattacharjee and Liu [14] use ChatGPT to detect LLM-generated text. There are also works focused on designing prompts to enable LLMs to detect suspiciousness content such as clickbait [291], hate speech [29], and sexual predatory [156].

5.1.4 Scoring Function

LLMs can also be used as scoring or ranking functions that rate and rank items based on the user’s behavior history. Chat-REC [63] and Liu et al. [133] leverage frozen LLMs to conduct item rating prediction. Kang et al. [100] comprehensively evaluate LLMs item rating prediction in zero-shot, few-shot, and fine-tuning scenarios. Dai et al. [36] formulate point-wise, pair-wise, and list-wise to elicit LLM’s item scoring ability. BookGPT [288] employs LLMs to predict both book ratings and personalized user ratings on books. TallRec [8] employs instruction tuning for book and movie rating prediction. RecLLM [61] takes LLMs as a ranking and explanation module by feeding side information to produce a score for an item and an explanation for the score. Hu et al. [78] use LLMs to predict the user’s satisfaction score with a conversation response and select the response with the highest satisfaction as output.

5.1.5 Explainer

The strong generation and reasoning ability make LLMs good explainers for the UM system, making its process understandable to humans. Liu et al. [133] prompts LLMs to generate recommendation explanations based on the user’s interaction history. Chat-REC [63] explains the recommendation in the interaction process with users. Yu et al. [265] fine-tune Open-Llama to generate an explanation of the financial time series data forest. ALEX [97] uses LLM to generate an explanation for health-related UGC classification and would self-correct the prediction after explaining. Bao et al. [9] propose a self-reinforcement LLM-based framework in learnersourcing to generate student-aligned explanations, evaluate, and iteratively enhance explanations automatically. Wang et al. [227] use LLMs to generate explanations for hateful and non-hateful content. Ziems et al. [292] investigate LLMs capabilities in social content explanation generation tasks like figurative language explanation and implied misinformation explanation.

5.1.6 Chatbot

LLMs are often used as chatbots to empower the UM system’s interactivity and enhance understanding of UGC. Chat-REC [63] employs ChatGPT as a ChatBot to interact with human by incorporating user queries with user profiles and user behavior history. Lin and Zhang [129] observe ChatGPT’s behavior in recommendation-oriented dialogues and demonstrate the potential for ChatGPT to serve as an artificial general recommender (AGR). [178] use ChatGPT as a chat interface with humans to understand the humans’ mental status based on their tweets and responses. Lakkaraju et al. [112] investigate LLMs’ potential in serving as personal financial advisors, where LLMs iteratively interact with humans as a chatbot. Hassan et al. [68] present a framework that enables LLMs to act as personal data scientists, where users interact with an LLM-based chatbot, and LLMs would return data analysis. CharacterChat [217] designs prompts with behavior preset and dynamic memory to help LLMs act as a chatbot with a specific personality. Zheng et al. [286] augment Llama with ChatGPT-generated dataset for optimized emotional support chatbot. He et al. [72] empirically study conversational recommendation with LLMs by constructing in-the-wild datasets. ChatDoctor [127] fine-tunes Llama using 100,000 patient-doctor dialogues to create a specialized Chatbot with enhanced accuracy in medical advice. Chen et al. [23] propose a three-stage pipeline to design prompts for ChatGPT to serve as both doctor and patient chatbots. GeneRec [233]

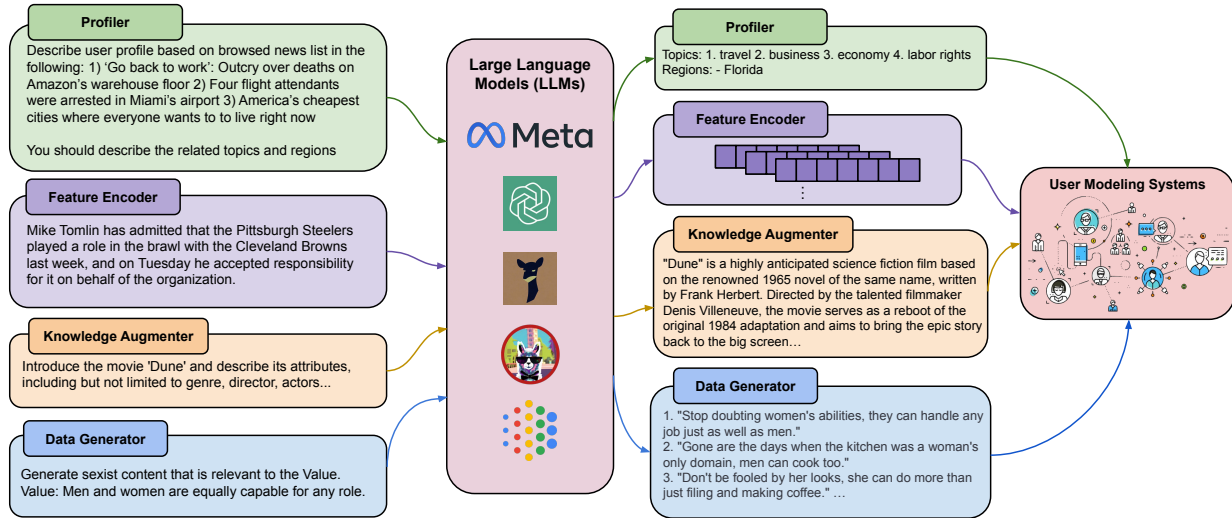


Figure 14: LLMs-as-Enhancers, where LLMs are leveraged to generate user profiles, content embeddings, knowledge-augmented content, and training data to augment downstream user modeling systems.

takes ChatGPT as a user conversational interface and takes user instruction and feedback to generate personalized content.

5.2 LLMs-as-Enhancers

In this section, we analyze the approaches that leverage LLMs as enhancers in the UM models. That being said, LLMs are not used to generate task answers directly but are leveraged to serve as plug-in augmentation modules instead. Approaches use LLMs as profilers, feature encoders, knowledge augmenters, and data generators.

5.2.1 Profiler

Using LLMs-as-profilers involves the creation of prompts based on users' history, including their watching, purchasing, and viewing activities. These prompts are then inputted into LLMs to generate various aspects of users' profiles, such as their characteristics, personality, geographical location, and areas of interest. The resulting user profiles are represented in natural language, making them easily understandable to humans. They are commonly employed as input for tasks such as recommendation and rating prediction, enabling downstream predictions to be tailored to the specific requirements of individual users. HKFR [260] utilizes user heterogeneous behavior, encompassing behavior subjects, behavior content, and behavior scenarios, and feeds it into ChatGPT to obtain user profiles. ONCE [136] and GENRE [135] employ LLMs to generate topics and regions of interest based on user browsing history. PALR [26] uses an LLM and user behavior as input to generate user profile keywords. KAR [246] leverages LLMs to generate user and item profiles, encompassing user preferences and factual knowledge about items, respectively. LGIR [46] completes users' resumes by incorporating explicit properties from their self-description and implicit characteristics from their behavior history. GIRL [287] leverages LLMs to generate suitable job descriptions based on the user's curriculum vitae to help the recommendation model better understand the job seeker's preferences. NIR [229] feeds user-watching history and design prompts for LLMs to generate user preference to augment recommendations. Once these user profiles, which include information such as age, gender, preferences, topics of interest, and geographic location, are obtained, they can be fed into the downstream user modeling system to enhance understanding of user preferences.

5.2.2 Feature Encoder

Given that LLMs have incredible user-generated content understanding and modeling capabilities, some research focuses on using LLM-generated text embeddings to enhance UM systems. GPT4SM [171] uses both PLMs and GPT to encode recommendation queries and candidate text for relevance prediction. LKPNR [187] uses open-source LLMs such as Llama and RWKV [169] to get news representation for better semantic capture capability. Li et al. [123] explore LLMs' potential in text-based collaborative filtering by using LLMs with parameters ranging from 125M to 175B as feature encoders. LLM4Jobs [121] leverages LLMs to embed both job posts and occupation taxonomy database and calculate their embedding similarity for recommendations. SentimentGPT [104] uses GPT to encode text embeddings for small ML model training on sentiment analysis tasks. KAR [246] leverages ChatGLM as a knowledge encoder to get latent embeddings for user profiles and item factual knowledge. These LLM-based text embeddings are then fed into the downstream models to inject the LLMs' UGC understanding ability into downstream task-agnostic models. LLMs are shown to have strong natural language understanding ability and rich open-world knowledge. Armed with LLM-encoded embeddings, the downstream UM systems could better understand the semantic information in UGC and leverage the world knowledge from LLM latent space.

5.2.3 Knowledge Augmenter

LLMs have demonstrated impressive capabilities in internalizing knowledge and responding to common queries [159, 161], which can serve as knowledge bases and bring factual knowledge to UM systems. Yin et al. [260] leverages LLMs to fuse diverse heterogeneous knowledge, including multiple behavior subjects, multiple behavior contents, and multiple behavior scenarios. Mysore et al. [154] augment narrative-driven recommendation using LLMs for author narrative query generation based on user-item interactions and train retrieval models with these LLM-augmented queries. KAR [246] prompts LLMs to generate factual knowledge relevant to the item for recommender system augmentation. GPT4Rec [120] leverages GPT to expand users' search queries to give item titles and users' history and feed item titles for rephrasing. Li et al. [122] use prompt tuning in LLMs to generate aspect embedding extraction and then feed the embeddings into aspect-based recommendation systems. LLM-Rec [145] proposes four prompting strategies to encourage LLMs to generate knowledge-augmented item descriptions for recommendation. Acharya et al. [3] leverage LLMs to generate detailed item descriptions for knowledge-augmented recommendations. Fang et al. [52] leverage ChatGPT knowledge to generate rephrases of training datasets instances to enhance model generalization and performance on unseen compositions. PULSAR [118] integrates LLMs in data augmentation to generate more knowledge relevant to the annotated training data. Knowledge stored in LLMs can also be used to augment dialogue. AugESC [284] uses LLMs to complete full dialogue and construct a scalable dataset to augment dialog systems' generalization ability in open-domain topics. Schlegel et al. [196] ask LLMs to generate a hypothetical conversation between the doctor and the patient based on a medical note and then use this data to train a specialized LM. Cohen et al. [35] leverage GPT-3 to do back translation and rephrase as data augmentations for hate speech detection. TacoBot [150] leverages LLMs to synthesize user intents of conversation for data augmentation. LLM-PTM [268] proposes an LLM-based privacy-aware augmentation for patient-trial matching, which leverages LLMs to create supplementary data points while preserving the semantic coherence of the original trail's inclusion and exclusion criteria. The collective findings of these works show that researchers could enhance UM systems performance by identifying and injecting external reasoning or factual knowledge into UM text input. Take movie recommendation as an example, LLMs provide contextual knowledge for candidate item descriptions under specific scenarios. This approach paves the way for the advancement of open-world recommendation, which integrates broader contextual information to enhance the recommender system.

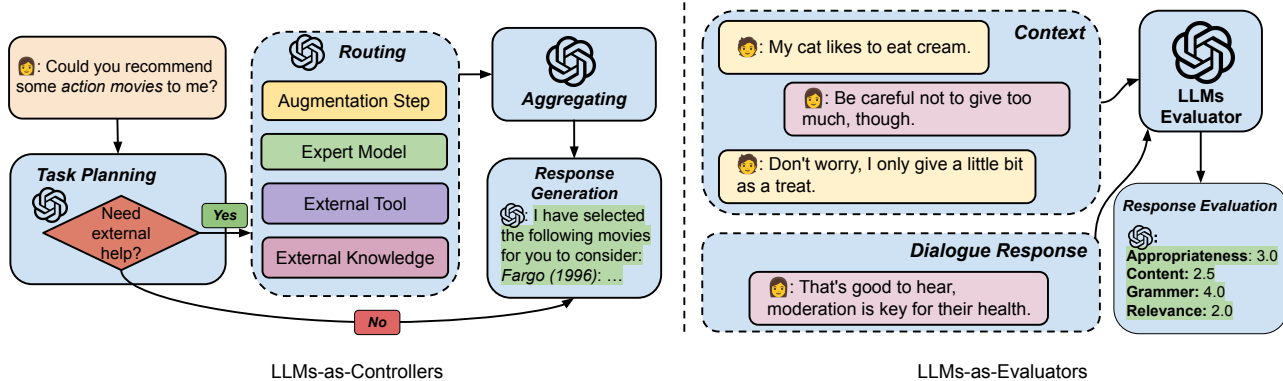


Figure 15: LLMs-as-Controllers and LLMs-as-Evaluators, where LLMs are employed to manage the LLM-UM pipeline and to evaluate the output of UM systems, respectively.

5.2.4 Data Generator

Owe to LLMs’ strong sequence generation and data compression ability [267], there is a lot of research that leverages LLMs to generate training data or weak labels and feed these into small UM models for training [16, 170, 222], which can also be envisioned as knowledge distillation. VA-Models [7] generates value-aligned training data from LLMs by few-shot prompting, then uses the generated data to train small models on value-alignment judgment tasks. Tang et al. [211] investigates LLMs’ potential to synthesize high-quality clinical data with labels and fine-tune local models for the downstream tasks. LLM4Job [121] leverages GPT-4 to generate occupation coding datasets. Su et al. [203] investigate fake news detection models’ bias with LLMs-synthesized fake news articles. Veselovsky et al. [221] find LLMs-synthesized data have a different distribution than real-world data, and propose three strategies: ground, filter, and taxonomy-based generation strategy to combat this difference, which has proved to be effective in sarcasm detection. LLMs can also generate weak labels to enhance UM systems. For example, Leite et al. [114] propose to prompt LLMs with credibility signals to produce weak labels to enhance misinformation detection performance. Deng et al. [41] propose to generate weak financial sentiment labels for Reddit posts with an LLM and then use that data to train a small model that can be served in production. Foosherian et al. [60] find that LLMs can aid conversational agents in generating training data, extracting entities and synonyms, localization, and persona design. For the LLM-generated text detection domain, LLMs are employed to synthesize training corpus [21, 138, 262, 266]. Graph-Toolformer [273] uses ChatGPT to annotate and augment a large prompt dataset that contains API calls of external reasoning tools and uses the synthesized dataset to fine-tune open-source LLMs. RefGPT [252] uses LLMs to generate truthful and customized dialogues without hallucination.

LLMs can also generate high-quality conversational data. Zheng et al. [286] recursively prompting ChatGPT with in-context learning to generate an extensive emotional support dialogue dataset (ExTES) and use it to fine-tune Llama for optimized emotional support dialogue systems. Wang et al. [228] propose a LLMs cooperation system named a doctor-patient loop to generate high-quality conversation datasets. Overall, these works indicate that by training on LLM synthesized data, small UM models could inherit the strong user understanding ability from LLMs. These approaches are advantageous in low-resource circumstances or some applications that need efficient deployment. For instance, in healthcare applications where privacy concerns and data scarcity are prevalent, utilizing LLMs to generate supplementary data can significantly alleviate the low-resource and privacy problems and thus bolster the effectiveness of UM systems.

5.3 LLMs-as-Controllers

The scale of parameters of LLMs brings emergent abilities that have never been observed in small language models and gives LLMs unprecedented ability to control the system pipeline and supercharge the UM system for personal needs. Note that different from LLMs as agents that let LLMs freely explore and interact with the environment, LLMs-as-Controllers include works that have designed the entire pipeline and let LLMs decide whether to conduct certain operations. HuggingGPT [198] employs LLMs as a controller to manage and organize the cooperation of expert models. Specifically in the user modeling systems, RecLLM [61] leverages LLMs as a dialogue manager, which converses with the user, tracks context, and makes system calls when necessary. Chat-REC [63] lets LLMs decide when to use recommendation systems as external tools. Foosherian et al. [60] demonstrate that LLMs can assist the pipeline-based conversational agent in contextualization, intent classification to prevent conversational breakdown and handle out-of-scope questions, auto-correcting utterances, rephrasing responses, formulating disambiguation questions, summarization, and enabling closed question-answering capabilities. LLM4Jobs [121] constructs a pipeline that makes LLMs determine whether to do a summary on the job posting.

5.4 LLMs-as-Evaluators

Evaluating natural language generation (NLG), especially in the open-domain and conversation settings, has posed significant challenges in user modeling. The strong language modeling capabilities in LLMs open up new opportunities for these complex evaluations, and some research works propose to envision LLMs as evaluators for generative UM systems. Svikhnushina and Pu [207] leverage LLMs to approximate online human evaluation for dialogue systems. Huynh et al. [84] explores LLMs capabilities in dialog evaluation and the relation between prompts and training datasets. iEvalM [235] proposed a conversational recommendation evaluation framework that leverages LLMs to simulate various interactions between users and systems. Zheng et al. [285] discover that using strong LLMs like GPT-4 as a judge can match both controlled and crowdsourced human preferences well. LLM-Eval [132] designs a prompt-based evaluation method that leverages a unified evaluation schema to cover multiple dimensions of conversation quality. Bhat et al. [13] takes a fine-tuned GPT-3 to evaluate the generated questions by classifying if the questions are useful to learning. GIRL [287] evaluates the recommended job results with the help of ChatGPT. These works reveal that LLMs can be an effective tool of assessing UM system output, determining the extent these outputs are customized to meet specific user needs. Particularly in conversational contexts, where conducting user studies could be expensive and prone to bias, LLMs provide a reliable and effective approach to assess the quality of complex and open-ended generations. Consequently, the LLMs-as-Evaluator paradigm enhances UM system development.

Personalized Prompting.

It is worth mentioning that to make LLM-UM adapt to an individual’s specific needs and generate personalized content, existing research works design prompting templates to encode UGC and user behavior history to help LLMs understand users’ preferences. Existing personalized prompting paradigms fall into three categories: vanilla, retrieval-augmented, and profile-augmented methods. We provide specific examples of these prompts in the context of the personalized product rating task below.

- **Vanilla Personalized Prompt:** Here is the user rating history: {{all user behavior history}}. Based on the above rating history, please predict the user’s rating for the product: {{query item}}.
- **Retrieval-Augmented Personalized Prompt:** Here is the user rating history: {{top-k user behavior history}}. Based on the above rating history, please predict the user’s rating for the product: {{query item}}.

- **Profile-Augmented Personalized Prompt:** `{{User Profile}}`. Here is the user rating history: `{{top-k user behavior history}}`. Based on the above rating history, please predict the user’s rating for the product: `{{query item}}`.

Most works under the LLMs-as-Predictors paradigm adopt the in-context learning paradigm and encode the entire user behavior history as in-context examples (**Vanilla Personalized Prompt**). BookGPT [288], for instance, employs a few-shot prompting strategy to enable LLMs to understand the correlation between book content and personalized ratings. Dai et al. [36] encode user rating history as few-shot demonstration examples. Liu et al. [133] supply the LLM with the user’s interaction history relating to the task, enabling the generation of personalized content. Christakopoulou et al. [32] utilize few-shot prompting to demonstrate the user research journey to LLMs. and find that feeding partial user behavior log causes significantly lower performance. Moreover, BookGPT [288] discovers that a longer user behavior history would bring better performance. Considering the increasing length of UGC and user behavior history and the limited context length of LLM, some studies have applied retrieval methods to select the most relevant user data for enhancing LLM personalization (**Retrieval-Augmented Personalized Prompt**). For instance, LaMP [189] introduces a retrieval-augmented method to obtain the most relevant content in behavioral history and incorporate it into the prompt.

Except for previously mentioned works that fall into the LLMs-as-Predictors paradigm, some works can be categorized into the LLMs-as-Enhancers paradigm that employs LLMs to rephrase and summarize user preferences and profiles based on their history (**Profile-Augmented Personalized Prompt**). Richardson et al. [185] suggest the use of instruction-tuned LLMs to generate an abstract summary of user data, augmenting retrieval-based personalized methods like LaMP. HKFP [260] inputs user behavior history into LLMs to fuse heterogeneous user knowledge, which assists fine-tuned LLMs in understanding user preferences. Note that the generated profile could vary in different tasks. For example, in rating prediction tasks, Richardson et al. [185] prompt instruction-tuned LLMs to summarize the user’s most common positive and negative score as profile, while generating user research interest as profile in the personalized citation identification task.

6 Applications of LLM-UM

We introduce the applications that can be categorized into personalization and suspiciousness detection.

6.1 Personalization

Personalization in user modeling refers to tailoring and customizing a system’s interactions, content, or recommendations to meet the specific needs, preferences, and characteristics of individual users. In this section, we look into LLM-UM for personalization applications, including user-generated content (UGC) analysis, user behavior prediction, personalized assistance, personalized recommendation, personalized dialog system, personalized education, and personalized healthcare.

6.1.1 User Profiling

User profiling refers to mining the characteristics, personality, and potential preferences based on UGC and user behavior, paving the way for downstream personalized applications. User profiling mainly includes detecting users’ stances and sentiments and analyzing users’ characteristics, personalities, etc. Zhang et al. [271] discuss LLMs’ impact in stance detection topics and the opportunities they bring. LoT [76] is proposed to help LLMs assimilate high-quality external knowledge to boost stance detection. Mu et al. [153] and Ziems et al. [292] comprehensively investigate LLMs in the computational social science tasks, including detecting sarcasm, hate speech, ideology, stance, and more in UGC. SentimentGPT [104] is a pioneering to leverage LLM for UGC sentiment analysis. Wu et al. [244] discover that LLMs can be used to estimate the latent stance of politicians

and give solutions to complex social science measurement problems. For general personality modeling, Rao et al. [184] use LLMs to assess human personality based on the MBTI test. Ji et al. [87] employ a level-oriented prompting strategy to analyze the user’s personality in a given text. Christakopoulou et al. [32] use LLMs to mine users’ interest journey and provide deeper, more interpretable, and controllable user understandings.

6.1.2 Personalized Recommendation

LLMs are adopted in the recommendation system to offer personalized suggestions towards candidate items tailored to meet user preferences and specific needs. The recommendation tasks can be further categorized into top-k recommendation, rating prediction, and conversational recommendation.

Top-k Recommendation. The top-k recommendation task directly predicts the top-k favorite items based on the user’s behavior history. Most methods directly design prompts feeding user behavior history and optional characteristics into LLMs to generate recommended items. In context learning is the most common paradigm for the top-k recommendation, which gives several exemplars and recommendation results to help LLMs better understand tasks. Representative works include [36, 46, 73, 133, 135, 235, 274]. Zhang et al. [276] further utilized Chain-of-Thought prompting to conduct a top-k recommendation task. For better representation and domain adaptation, researchers also fine-tune LLM’s parameters. Chen [26] and GenRec [86] fine-tune a Llama-7B model to help the model adapt to recommender system and generate items. Zhang et al. [276] conduct instruction tuning on a Flan-T5-XL model to help it adapt to recommendation. GIRL [287] further proposes a reward model and uses reinforcement learning to provide better feedback for LLMs fine-tuning.

Rating Prediction. The rating prediction means predicting the user’s rating for given items. In this process, LLMs give the predicted scores for items in the context of the user’s behavior and UGC history. The rating prediction task can probe LLM’s capabilities in understanding user preference and can also be understood as user behavior prediction. Similar to top-k recommendation, rating prediction can also be categorized into frozen LLMs and fine-tuning LLMs. Armed with frozen LLMs, BookGPT [288] and Dai et al. [36] conduct prompt engineering to make LLMs generate predicted ratings; KAR [246] takes LLMs to generate user profile and factual knowledge of items and feed them into the discriminative recommendation for rating prediction. For rating prediction with fine-tuned LLMs, Kang et al. [100] fine-tune LLMs in the rating prediction task; TallRec [8] integrate LoRA and instruction tuning for rating prediction task adaptation; GLRec [242] incorporates meta-paths into soft prompt and conduct instruction tuning for item generation. Further, Graph-Toolformer [273] fine-tunes LLMs and empowers them to use external graph reasoning tools for user rating prediction.

Conversational Recommendation. Conversational recommendation means using the system to interact with the user and understand the user’s preference through conversation and then generate recommended items in the conversation. Chat-REC [63] converts the user profile and historical interactions into prompts to build conversational recommendation systems. GeneRec [233] adopts ChatGPT to personalize content generation, and it leverages user instructions to acquire users’ information needs. Lin and Zhang [129] envision ChatGPT as an Artificial General Recommender (AGR) that comprises conversationally and universality to engage in natural dialogues and generate recommendations across various domains.

6.1.3 Personalized Assistance

Personalized assistance refers to using LLM techniques to tailor and customize generated content based on individual preferences, behavior, and characteristics of users. Tian et al. [214] empirically study the LLMs’ potential as fully automated programming assistants in the tasks of code generation, program repair, and code summarization. TacoBot [150] is an LLM-augmented task-oriented dialogue system that guides users through complex real-world tasks with multiple steps. LaMP [189] uses LLMs to conduct personalized text classification and personalized text generation, such as personalized citation identification and personalized news headline

generation, etc. DISC-LawLLM [269] is an intelligent system that utilizes LLMs to provide a wide range of personalized legal advice. FinGPT [142] is an LLM that can offer personalized investment advice based on the user's risk and financial goals. Chakrabarty et al. [20] investigate LLMs' ability in creative writing assistance, including planning, translating, and reviewing processes.

6.1.4 Personalized Dialogue System

LLMs are also widely adopted in dialogue systems and combined with the user's behavior history and preference to provide personalized user experiences. Hudeček and Dušek [83] utilized LLMs to retrieve the user behavior context and user history for personalized conversational response generation. DiagGPT [19] extends LLMs to task-oriented dialogue scenarios, in which LLMs need to pose questions and guide users toward specific task completion proactively. Cho et al. [30] use GPT-2 to generate dialogue data while detecting the user's persona. RefGPT [252] proposes to generate enormous truthful and customized dialogues without worrying about factual errors caused by the model hallucination to enable the personalized dialogue system.

6.1.5 Personalized Education

LLMs have shown great potential in promoting the equality of education and improving the existing education paradigm by adapting LLM-based tools to tailor for students and instructors [98, 249], and education for data scientists [218]. Koyuturk et al. [109] discover that with multiple interaction turns, LLMs can adapt the educational activity to the user's characteristics, such as culture, age, and level of education, and its ability to use diverse educational strategies and conversational styles. EduChat [37] is an LLM-based chatbot that aims to strengthen personalized, fair, and compassionate intelligent education, serving teachers, students, and parents. Sharma et al. [197] investigate ChatGPT's performance on the United States Medical Licensing Examination (USMLE) and point out that ChatGPT can be an invaluable tool for e-learners. Elkins et al. [48] investigate ChatGPT's performance in generating educational questions in classroom settings and find them high-quality and sufficiently useful. Ochieng [157] delve into LLM's ability to participate in educational guided reading, including generating questions based on the input text and recommending content based on the user's response. C-LLM [158] examines the implications for LLMs for AI review and assessment of complex student work. Phung et al. [174] comprehensively investigate ChatGPT's capability in a set of programming education scenarios and find GPT-4 comes close to human tutor in several scenarios.

6.1.6 Personalized Healthcare

LLMs also play an important role in empowering healthcare services and providing personalized service. Liu et al. [139] discover that LLMs are capable of grounding various physiological and behavioral time-series data and making meaningful inferences under the few-shot settings. Wang et al. [237] investigates the performance of LLMs on clinical language understanding tasks and introduces self-questioning prompting to enhance LLM in clinical scenarios. HeLM [10] enables LLMs to use high-dimensional clinical modalities to estimate underlying disease risk with individual-specific data. PharmacyGPT [143] utilizes LLMs to generate comprehensible patient clusters, formulate medication plans, and forecast patient outcomes. Zhang et al. [275] utilize LLMs to identify patients with specific medical diagnoses and provide diagnostic assistance to healthcare workers. Zhongjing [255] introduces a Llama-based LLM with expertise in the Chinese medical domain and can provide personalized advice for the user's specific case.

LLMs are widely applied for mental health. Ghanadian et al. [64] utilize LLMs for suicidality assessment from social media. Qi et al. [177] takes LLMs to evaluate suicidal risk and cognitive distortion identification on Chinese social media platforms. Wang et al. [228] leverage LLMs to generate note-oriented doctor-patient conversations. Chen et al. [23] utilize ChatGPT to simulate conversations between psychiatrist and patient based

on user experience. Fu et al. [62] present an LLM-empowered framework that assists non-professionals in providing psychological interventions on online user discourse. Mental-LLM [248] investigate LLM’s capabilities in mental health tasks and find the superiority of instruction tuning in boosting LLMs’ performance in mental health tasks. Lai et al. [111] propose an LLM-based assistant for question-answering in psychological consultation settings to ease the demand for mental health professions. Peters and Matz [173] assess the ability of GPT-3.5 and GPT-4 to infer the psychological dispositions of individuals from their digital footprints.

6.2 Suspiciousness Detection

User modeling involves the process of understanding and predicting users’ behavior and preferences. By creating a comprehensive model of user’s normal behaviors, it becomes possible to detect deviations from this norm. For users with suspicious behavior history, UM system could potentially isolate or warn users who exhibit harmful behaviors, while offering a more open environment to those with positive engagement histories. Therefore, suspiciousness detection is a key application of user modeling. Suspiciousness detection refers to the process of identifying or recognizing behaviors, actions, or patterns that are deemed to be suspicious or potentially indicative of anomalous, illegal, harmful, or malicious activities [92, 93, 281]. This section introduces leveraging LLM-UM to detect fraud, hate speech, misinformation, misconduct, and LLM-generated content.

6.2.1 Fraud Detection

Suspiciousness in fraud detection refers to fraudulent and deceptive behavior, which is widely adopted in financial transactions, social networks, and more. Some research leverages LLMs to supercharge fraud detection models. [291] design prompts to enable LLMs for clickbait detection and achieves satisfied performance. Yang and Menczer [253] present a case study on a Twitter botnet that employs ChatGPT to generate human-like content, while state-of-the-art LLM content classifiers fail to detect them. Spam-T5 [110] fine-tunes T5 for spam email detection and outperforms baselines with limited training samples. Ayoobi et al. [6] leverage LLMs to generate fake LinkedIn profiles and develop the Section Tag Embeddings to detect fake profiles. Shukla et al. [200] explores GPT-3 and GPT-4 for fraudulent physician review detection. Ziems et al. [293] employs GPT-4 to explain the decisions of classical machine learning models on network intrusion detection.

6.2.2 Discrimination Detection

LLMs’ strong natural language understanding ability can supercharge hate speech detection, which is of great importance in social content moderation [147]. Chiu et al. [29] leverage GPT-3 to identify sexist and racist UGC under zero-shot, one-shot, and few-shot settings. Cohen et al. [35] leverage GPT-3 to back translate and rephrase as augmentations for hate speech detection. Del Arco et al. [40] explores using LLMs with prompting for zero-shot hate speech detection. Das et al. [39] evaluate LLMs’ capabilities in multilingual and emoji-based hate speech. Wang et al. [227] further prompt LLMs to generate an explanation for hateful and non-hateful content and investigate the explanation generated by LLMs. Nguyen et al. [156] fine-tune Llama-7B to detect online sexual predatory chats and abusive language.

6.2.3 Misinformation Detection

LLMs can also be leveraged to detect misinformation, especially fake news [89]. Chen and Shu [21] discover that LLM-generated misinformation can be harder to detect compared to human-written, which suggests a more deceptive style and potentially causes more harm. Pavlyshenko [167] explores using a fine-tuned Llama-2 model for misinformation analysis and fake news detection. Yang et al. [251] assess ChatGPT’s ability to rate the credibility of news outlets and find ChatGPT’s prediction correlates to those from human experts. Pan et al. [164] establish a threat model and reveal that LLMs can act as effective misinformation generators, leading to

significant degradation in open-domain question-answering systems. Leite et al. [114] develop a misinformation detection approach that combines the zero-shot LLM credibility signal labeling and weak supervision and achieve state-of-the-art without using ground truth label for training. Su et al. [203] discover that existing detectors are prone to flagging LLM-generated text as fake news and propose to leverage adversarial training for bias mitigation. Huang and Sun [82] explore ChatGPT’s proficiency in generating, explaining, and detecting fake news, and propose using potential extra information that could boost LLM-based fake news detection.

6.2.4 LLM-Generated Text Detection

As LLMs emerge, misuse of LLMs increases, including disseminating fake news, plagiarism, manipulating public opinion Hanley and Durumeric [66], cheating, and fraud, making detection of LLM-generated content essential [21, 43, 144, 168, 210]. Researchers find that LLM-generated text can easily evade the plagiarism-checking software Khalil and Er [103] and is difficult to be identified by humans [223]. SpaceInfi [18] discovers that the extra space can serve an important role for LLM-generated content to evade detection.

GPT-Pat [266] develops a framework consisting of a Siamese network to compute the similarity between original and LLM-generated text and a binary classifier. Yang et al. [254] measure the ChatGPT’s involvement in text generation based on edit distance. Orenstrakh et al. [160] investigate the LLM-generated text detection tool specifically in the education domain. CoCo [138] presents a coherence-based contrastive learning method to detect LLM-generated text under the low-resource scenario. Bhattacharjee and Liu [14] employ ChatGPT as a detector for AI-generated text detection.

7 Challenges and Directions

In this survey, we have comprehensively reviewed the recent approaches and applications of LLM-UM. Since the integration of LLMs into user modeling systems is still in the early stage, there are still some important and challenging yet unsolved problems in this direction. In this section, we discuss some challenges and potential future directions in this field.

7.1 Hallucination Mitigation

Though LLMs have shown strong capabilities in various domains, a significant challenge is the hallucination problem [88, 279], where LLMs generate seemingly plausible content but deviate from user input [4], previously generated content [137], and factuality knowledge [56, 149]. In the context of user modeling, the challenges of hallucination can be categorized into 1) *factual accuracy*: LLMs can sometimes generate content that sounds plausible but is either factually incorrect or not grounded in the input data. 2) *hallucination in high-stakes scenarios*: In contexts where recommendations have profound implications, such as in healthcare or legal advice, hallucinations can have severe consequences. 3) *user intent misunderstanding*: Hallucination can also arise from a misunderstanding of user intent and input context, which can significantly hinder the performance of LLMs in user modeling.

To mitigate the prevalent hallucination problems in LLMs, a promising approach to counter the hallucination problem in LLM-UM is to incorporate trustworthy symbolic knowledge in the LLMs’ input, including unifying knowledge graphs (KGs) [163] and the use of external knowledge retrieval for augmented generation [264]. Manipulating the output side of LLMs can also help alleviate the hallucination problem. This could involve generating calibrated confidence scores to validate predictions and encouraging LLMs to express uncertainty. Research could also focus on designing decoding strategies that foster more reliable content generation, including factuality feedback mechanisms for the next token selection.

7.2 Privacy and Security

Privacy and security are paramount concerns in the user modeling system since it can involve processing and analyzing a wealth of UGC and user behavior history. However, recent research suggests that LLMs can pose privacy risks. For example, an adversary can recover training examples containing a person’s name, email address, and phone number by querying the model [79]. The challenges in LLM-UM can be summarized as 1) *balancing personalization and privacy*: The success of LLMs in user modeling relies heavily on personalization, which necessitates the use of sensitive user data. However, this can lead to privacy and security concerns. That being said, striking the right balance between personalization and privacy is a significant challenge. 2) *data leakage risk*: As LLMs are trained on large quantities of data, there’s a risk that they might unintentionally leak sensitive information during the text generation process. Recent research has shown that LLMs are vulnerable to adversarial attacks by prompt injection [128] and jailbreaking [183]. This could potentially leak users’ identities or private information and cause security issues. 3) *adaptation to evolving privacy laws*: Privacy laws and regulations are continually evolving, and LLMs need to be updated to ensure compliance, which can be a challenging task given the huge parameters and high updating cost.

To mitigate the above challenges, future research can focus on designing privacy techniques in preprocessing, training, and post-processing steps to mitigate user privacy risks. They can study how to provide users with more transparency and study how their data are used and give them more control over their information.

7.3 Complex Structural Data Understanding

User interactions fundamentally form complex, heterogeneous, temporal text-rich graphs, encapsulating user-user and user-item interactions. Therefore, equipping LLMs with the capability to comprehend and interpret such graph structures could dramatically improve the performance of user modeling. Current challenges of leveraging LLMs to model complex graph structure can be summarized as 1) *lack of sequential transformation*: Unlike natural languages, graph-structured data often lack a straightforward transformation into sequential text [224], which makes it difficult for LLMs to process and understand. 2) *gap between graphs and human languages*: The significant gap between graph structures and natural language used for LLMs pretraining can hinder the application of LLMs to facilitate graph reasoning. 3) *graph variability*: Graphs can define structure and features in distinct ways and lack a unified paradigm, making it hard to generalize well to the diverse structure and feature representations of different graphs [96, 225]. 4) *heterogeneous temporal text-rich graph structure modeling*: User interactions can be formulated into heterogeneous temporal text-rich graphs, posing challenges for LLMs to understand such complex data.

To empower LLMs with the capabilities to understand complex structural data, future research could focus on 1) *graph language generation*: One potential research direction could be to derive a language for graph structures that LLMs can better understand, such as transforming graph data into a form that LLMs can process, such as sequential text. 2) *temporal heterogeneous graph modeling*: Given that user interactions often form heterogeneous temporal graphs, the research could focus on techniques to handle such complex structures, which could involve developing methods to capture the dynamics and heterogeneity of user interactions.

7.4 Comprehensive Benchmark and Evaluation Criteria

Evaluation is paramount to the success of LLM-UM, which can help us better understand the strengths and weaknesses, ensuring safety and reliability, and inspiring future research directions of user modeling systems. Current challenges of LLM-UM evaluations lie in 1) *lack of AGI-level benchmark*: Current evaluation benchmarks mainly focus on restricting models input and output to a static pair based on human annotation, which cannot measure the superhuman capabilities in LLMs, instead, the existing benchmarks can be seen in LLMs pertaining stage and cause test data leakage problem. Therefore, designing more challenging tasks and more reliable evaluation mechanisms by utilizing the cross-disciplinary knowledge from education, social science, and more

could be a better solution to probe LLM-UM’s capabilities as a superhuman-level system. 2) *comprehensive intermediate evaluation*: A general concern in LLM-UM is the evaluation only focuses on the prediction performance on specific datasets. Although this is likely the most important metric, the intermediate steps such as user profile generation, LLM-generated factual knowledge and datasets, and pipeline control behavior are also important for researchers to more comprehensively understand the methods. 3) *dynamic and evolving evaluation*: Existing UM evaluation protocols rely on static and public datasets, which would lead to overfitting and could not assess the LLM-UM in evolving circumstances, especially in the context of fast-evolving user data. For example, AgentBench [140] is a pioneering work that designs evolving benchmarks and evaluates LLMs as agents, which is shown to be the better way of static input-output pair evaluation to evaluate LLM-based systems. Therefore, developing a dynamic and evolving testbed for LLM-UM to probe the system’s capabilities in adapting the evolving data is an important research direction.

7.5 Trustworthy User Modeling

LLM-UM can bring significant benefits to humans, including providing personalized advice, recommendations, assistance, and more. However, the black-box property and unreliable generation process of LLMs would pose a serious threat to users when the LLM-UM system is applied to high-stakes domains such as healthcare and finance. Therefore, how to enable trustworthy LLM-UM is of significant importance. Current challenges and potential future research directions include 1) *transparent and explainable reasoning*: Users need to understand how an LLM-UM system makes decisions to enhance trust. Therefore, promising research directions are how to explain the black-box LLMs and how to make the user modeling process transparent and explainable. 2) *reliability*: LLM-UM should provide consistent and accurate results, be robust to different input templates, and give consistent generation given the same or similar input to be trustworthy. 3) *user control*: Trustworthy LLM-UM systems should allow users to control how their data is used and how decisions are made, which could involve developing mechanisms for users to feedback, correct, and update the LLM-UM system behavior.

7.6 Fairness in LLM-UM

Fairness is a crucial aspect of user modeling systems, which secure unbiased user characteristics and profile generation and give fair assistance and recommendations for different groups of people. LLMs trained on extensive datasets can inadvertently learn and perpetuate human biases and stereotypes in training data, leading to unfair treatment or discriminatory predictions of specific user groups [57]. Moreover, balancing personalization and fairness is a significant problem, especially when user data reflects societal biases. Therefore, future LLM-UM research may include 1) *bias detection and measurement*: Detect bias that exists in LLM-UM systems training data, input, model, and task design, and develop fine-grained mechanisms to measure the bias from different aspects (e.g., political bias, ethnic groups bias, etc.) in LLM-UM systems. 2) *promote fairness in LLM-UM systems*: Mitigate bias and give non-discriminative predictions, including making LLM-UM more explainable and transparent to help better understand the source of unfair behavior and devise appropriate interventions. 3) *ethics and regulations*: Research should also focus on the ethical implications of LLMs and the potential for regulations to ensure fairness, which could involve interdisciplinary work, bringing together technologists, ethicists, sociologists, and lawmakers.

7.7 Efficient and Effective Domain Adaptation

LLMs are pretrained on a large general-purpose corpus and then plugged into user modeling systems. A lot of works use LLMs as a general-purpose reasoner with frozen pertaining parameters, while others fine-tune LLMs to help them adapt to the specific task in user modeling. The latter paradigm allows the model to leverage the general language understanding capabilities learned during the pretraining stage while specializing the knowledge

to the task at hand and shows superior task performance. However, fine-tuning and deployment can be highly computationally expensive, especially when deploying on millions and even billions of scale user data. Therefore, developing efficient domain adaptation methods to reduce the cost for LLMs to fix in user modeling systems and efficiently help LLM-UM be personalized and adapt to the user’s dynamic needs. Moreover, user-specific data is often sparse and noisy [38, 107], which can lead to overfitting during the fine-tuning process, which can limit the model’s ability to generalize to different users and contexts. Therefore, future research to combat these challenges can focus on 1) *efficient fine-tuning for LLM-UM*: There is a need to develop and refine techniques for more efficient fine-tuning of LLMs specific for user modeling context. Techniques such as sparsity induction, model pruning, and knowledge distillation could be explored to reduce the computational and memory footprints. 2) *online learning*: To cope with the dynamics of user interests, research could explore online learning strategies for LLM-UM systems. These strategies would continuously update the model based on the latest user-generated content and interactions.

7.8 Personalized LLM-UM Deployment

Recent research on personalized LLMs mainly focuses on two aspects: personalized prompt engineering and personalized preference alignment. For personalized prompt engineering works, researchers focus on designing prompts to help LLMs understand user preferences based on user behavior. As previously discussed, the personalized prompting paradigm would directly encode user behavior history, apply retrieval methods to obtain the most relevant user history, and leverage LLMs to summarize user preference to augment retrieval and generation performance. At the same time, the other line of work focuses on aligning LLMs with personal preferences using reinforcement learning. Personalized Soups [85] is a representative work that decomposes user preference into multiple dimensions and models the personalization into a multi-object reinforcement learning problem. However, the existing personalized prompting model exhibits limited generalization and context windows for encoding and comprehending user preferences from extensive user behavioral history. Furthermore, human alignment methods require explicit human preferences and pre-established preference dimensions, and the alignment finetuning process is computationally expensive. Therefore, future personalized LLM research should focus on 1) *parametric user preference mining*: As previous studies [141] have indicated, differential parameters are more effective in understanding and modeling user preferences compared to discrete natural language. Additionally, parametric personalization can infuse personalized knowledge more efficiently by using fewer tokens, thereby conserving the context length for LLMs. 2) *cross-model personalized knowledge transfer*: Personalized LLMs strive to diverge from the prevailing ‘one-size-fits-all’ paradigm, aiming to offer unique services tailored to each individual. However, given the ‘scaling law’ - which posits that larger models yield superior performance - not all users possess the computational resources to generate high-quality output by operating their own large models on local machines. Therefore, future research should work towards developing small models that can learn from user history and gain personalized knowledge on a local machine. These models should be able to transfer or plug their knowledge into a remote, large model to produce high-quality personalized output. 3) *parametric privacy preservation*: Recent findings suggest that encoded text embeddings alone can be used to reconstruct full text [151], posing potential privacy and safety risks. This becomes particularly relevant in personalization, where user preferences can be encoded using a small model on local machines. There exists a risk of private data reconstruction using latent embedding and parameters. Therefore, developing strategies to thwart such reconstruction while preserving latent embedding performance is a pressing issue.

8 Conclusion

Our work presents a comprehensive and structured survey of large language models for user modeling (LLM-UM). We show why LLMs are great tools for user modeling and understanding user-generated content (UGC) and user

interactions. We then review the existing LLM-UM research works and categorize their approaches that integrate text-based and graph-based UM techniques, including LLMs serving as enhancers, predictors, controllers, and evaluators. Next, we categorize the existing LLM-UM techniques based on their applications. Finally, we outline the remaining challenges and future directions in the LLM-UM field. This survey serves as a handbook for LLM-UM researchers and practitioners to study and use LLMs to augment user modeling systems, and inspires additional interest and work on this topic.

References

- [1] F. Abel, Q. Gao, G.-J. Houben, and K. Tao. Analyzing user modeling on twitter for personalized news recommendations. In User Modeling, Adaption and Personalization: 19th International Conference, UMAP 2011, Girona, Spain, July 11-15, 2011. Proceedings 19, pages 1–12. Springer, 2011.
- [2] F. Abel, Q. Gao, G.-J. Houben, and K. Tao. Twitter-based user modeling for news recommendations. In Twenty-Third International Joint Conference on Artificial Intelligence. Citeseer, 2013.
- [3] A. Acharya, B. Singh, and N. Onoe. Llm based generation of item-description for recommendation system. In Proceedings of the 17th ACM Conference on Recommender Systems, pages 1204–1207, 2023.
- [4] V. Adlakha, P. BehnamGhader, X. H. Lu, N. Meade, and S. Reddy. Evaluating correctness and faithfulness of instruction-following models for question answering. arXiv preprint arXiv:2307.16877, 2023.
- [5] R. Alghamdi and K. Alfalqi. A survey of topic modeling in text mining. International Journal of Advanced Computer Science and Applications, 6(1), 2015.
- [6] N. Ayoobi, S. Shahriar, and A. Mukherjee. The looming threat of fake and llm-generated linkedin profiles: Challenges and opportunities for detection and prevention. In Proceedings of the 34th ACM Conference on Hypertext and Social Media, pages 1–10, 2023.
- [7] Y. Bang, T. Yu, A. Madotto, Z. Lin, M. Diab, and P. Fung. Enabling classifiers to make judgements explicitly aligned with human values. arXiv preprint arXiv:2210.07652, 2022.
- [8] K. Bao, J. Zhang, Y. Zhang, W. Wang, F. Feng, and X. He. Tallrec: An effective and efficient tuning framework to align large language model with recommendation. arXiv preprint arXiv:2305.00447, 2023.
- [9] Q. Bao, J. Leinonen, A. Y. Peng, W. Zhong, T. Pistotti, A. Huang, P. Denny, M. Witbrock, and J. Liu. Exploring self-reinforcement for improving learnersourced multiple-choice question explanations with large language models. arXiv preprint arXiv:2309.10444, 2023.
- [10] A. Belyaeva, J. Cosentino, F. Hormozdiari, C. Y. McLean, and N. A. Furlotte. Multimodal llms for health grounded in individual-specific data. arXiv preprint arXiv:2307.09018, 2023.
- [11] S. Berkovsky, T. Kuflik, and F. Ricci. Entertainment personalization mechanism through cross-domain user modeling. In International Conference on Intelligent Technologies for Interactive Entertainment, pages 215–219. Springer, 2005.
- [12] M. Besta, N. Blach, A. Kubicek, R. Gerstenberger, L. Gianinazzi, J. Gajda, T. Lehmann, M. Podstawski, H. Niewiadomski, P. Nyczyk, et al. Graph of thoughts: Solving elaborate problems with large language models. arXiv preprint arXiv:2308.09687, 2023.
- [13] S. Bhat, H. A. Nguyen, S. Moore, J. Stamper, M. Sakr, and E. Nyberg. Towards automated generation and evaluation of questions in educational domains. In Proceedings of the 15th International Conference on Educational Data Mining, volume 701, 2022.
- [14] A. Bhattacharjee and H. Liu. Fighting fire with fire: Can chatgpt detect ai-generated text? arXiv preprint arXiv:2308.01284, 2023.
- [15] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. Journal of machine Learning research, 3(Jan): 993–1022, 2003.
- [16] V. Borisov, K. Seßler, T. Leemann, M. Pawelczyk, and G. Kasneci. Language models are realistic tabular data generators. arXiv preprint arXiv:2210.06280, 2022.
- [17] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell,

- et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [18] S. Cai and W. Cui. Evade chatgpt detectors via a single space. *arXiv preprint arXiv:2307.02599*, 2023.
- [19] L. Cao. Diaggpt: An llm-based chatbot with automatic topic management for task-oriented dialogue. *arXiv preprint arXiv:2308.08043*, 2023.
- [20] T. Chakrabarty, V. Padmakumar, F. Brahman, and S. Muresan. Creativity support in the age of large language models: An empirical study involving emerging writers. *arXiv preprint arXiv:2309.12570*, 2023.
- [21] C. Chen and K. Shu. Can llm-generated misinformation be detected? *arXiv preprint arXiv:2309.13788*, 2023.
- [22] J. Chen, Z. Liu, X. Huang, C. Wu, Q. Liu, G. Jiang, Y. Pu, Y. Lei, X. Chen, X. Wang, et al. When large language models meet personalization: Perspectives of challenges and opportunities. *arXiv preprint arXiv:2307.16376*, 2023.
- [23] S. Chen, M. Wu, K. Q. Zhu, K. Lan, Z. Zhang, and L. Cui. Llm-empowered chatbots for psychiatrist and patient simulation: Application and evaluation. *arXiv preprint arXiv:2305.13614*, 2023.
- [24] W. Chen, Y. Gu, Z. Ren, X. He, H. Xie, T. Guo, D. Yin, and Y. Zhang. Semi-supervised user profiling with heterogeneous graph attention networks. In *IJCAI*, volume 19, pages 2116–2122, 2019.
- [25] Y. Chen, R. Li, Z. Zhao, C. Peng, J. Wu, E. Hossain, and H. Zhang. Netgpt: A native-ai network architecture beyond provisioning personalized generative services. *arXiv preprint arXiv:2307.06148*, 2023.
- [26] Z. Chen. Palr: Personalization aware llms for recommendation. *arXiv preprint arXiv:2305.07622*, 2023.
- [27] Z. Chen, H. Mao, H. Li, W. Jin, H. Wen, X. Wei, S. Wang, D. Yin, W. Fan, H. Liu, et al. Exploring the potential of large language models (llms) in learning on graphs. *arXiv preprint arXiv:2307.03393*, 2023.
- [28] W.-L. Chiang, Z. Li, Z. Lin, Y. Sheng, Z. Wu, H. Zhang, L. Zheng, S. Zhuang, Y. Zhuang, J. E. Gonzalez, et al. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2023.
- [29] K.-L. Chiu, A. Collins, and R. Alexander. Detecting hate speech with gpt-3. *arXiv preprint arXiv:2103.12407*, 2021.
- [30] I. Cho, D. Wang, R. Takahashi, and H. Saito. A personalized dialogue generator with implicit user persona detection. *arXiv preprint arXiv:2204.07372*, 2022.
- [31] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
- [32] K. Christakopoulou, A. Lalama, C. Adams, I. Qu, Y. Amir, S. Chucuri, P. Vollucci, F. Soldo, D. Bseiso, S. Scodel, et al. Large language models for user interest journeys. *arXiv preprint arXiv:2305.15498*, 2023.
- [33] I. A. Christensen and S. Schiaffino. Entertainment recommender systems for group of users. *Expert systems with applications*, 38(11):14127–14135, 2011.
- [34] H. W. Chung, L. Hou, S. Longpre, B. Zoph, Y. Tay, W. Fedus, E. Li, X. Wang, M. Dehghani, S. Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022.
- [35] S. Cohen, D. Presil, O. Katz, O. Arbili, S. Messica, and L. Rokach. Enhancing social network hate detection using back translation and gpt-3 augmentations during training and test-time. *Information Fusion*, page 101887, 2023.
- [36] S. Dai, N. Shao, H. Zhao, W. Yu, Z. Si, C. Xu, Z. Sun, X. Zhang, and J. Xu. Uncovering chatgpt’s capabilities in recommender systems. *arXiv preprint arXiv:2305.02182*, 2023.
- [37] Y. Dan, Z. Lei, Y. Gu, Y. Li, J. Yin, J. Lin, L. Ye, Z. Tie, Y. Zhou, Y. Wang, et al. Educhat: A large-scale language model-based chatbot system for intelligent education. *arXiv preprint arXiv:2308.02773*, 2023.
- [38] A. Das, M. Degeling, D. Smullen, and N. Sadeh. Personalized privacy assistants for the internet of things: Providing users with notice and choice. *IEEE Pervasive Computing*, 17(3):35–46, 2018.
- [39] M. Das, S. K. Pandey, and A. Mukherjee. Evaluating chatgpt’s performance for multilingual and emoji-based hate speech detection. *arXiv preprint arXiv:2305.13276*, 2023.
- [40] F. M. P. Del Arco, D. Nozza, and D. Hovy. Respectful or toxic? using zero-shot learning with language models to detect hate speech. In *The 7th Workshop on Online Abuse and Harms (WOAH)*, pages 60–68, 2023.
- [41] X. Deng, V. Bashlovkina, F. Han, S. Baumgartner, and M. Bendersky. What do llms know about financial markets? a case study on reddit market sentiment analysis. In *Companion Proceedings of the ACM Web Conference 2023*, pages 107–110, 2023.
- [42] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer. Qlora: Efficient finetuning of quantized llms. *arXiv*

[preprint arXiv:2305.14314](#), 2023.

- [43] M. Dhaini, W. Poelman, and E. Erdogan. Detecting chatgpt: A survey of the state of detecting chatgpt-generated text. [arXiv preprint arXiv:2309.07689](#), 2023.
- [44] D. Di Palma, G. M. Biancofiore, V. W. Anelli, F. Narducci, T. Di Noia, and E. Di Sciascio. Evaluating chatgpt as a recommender system: A rigorous approach. [arXiv preprint arXiv:2309.03613](#), 2023.
- [45] Y. Dong, N. V. Chawla, and A. Swami. metapath2vec: Scalable representation learning for heterogeneous networks. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 135–144, 2017.
- [46] Y. Du, D. Luo, R. Yan, H. Liu, Y. Song, H. Zhu, and J. Zhang. Enhancing job recommendation through llm-based generative adversarial networks. [arXiv preprint arXiv:2307.10747](#), 2023.
- [47] S. T. Dumais. Latent semantic analysis. *Annual Review of Information Science and Technology (ARIST)*, 38: 189–230, 2004.
- [48] S. Elkins, E. Kochmar, I. Serban, and J. C. Cheung. How useful are educational questions generated by large language models? In *International Conference on Artificial Intelligence in Education*, pages 536–542. Springer, 2023.
- [49] W. Fan, Y. Ma, Q. Li, Y. He, E. Zhao, J. Tang, and D. Yin. Graph neural networks for social recommendation. In *The world wide web conference*, pages 417–426, 2019.
- [50] W. Fan, Z. Zhao, J. Li, Y. Liu, X. Mei, Y. Wang, J. Tang, and Q. Li. Recommender systems in the era of large language models (llms). [arXiv preprint arXiv:2307.02046](#), 2023.
- [51] Y. Fan and F. Jiang. Uncovering the potential of chatgpt for discourse analysis in dialogue: An empirical study. [arXiv preprint arXiv:2305.08391](#), 2023.
- [52] Y. Fang, X. Li, S. W. Thomas, and X. Zhu. Chatgpt as data augmentation for compositional generalization: A case study in open intent detection. [arXiv preprint arXiv:2308.13517](#), 2023.
- [53] M. Farid, R. Elgohary, I. Moawad, and M. Roushdy. User profiling approaches, modeling, and personalization. In *Proceedings of the 11th International Conference on Informatics & Systems (INFOS 2018)*, 2018.
- [54] S. Feng, Z. Tan, R. Li, and M. Luo. Heterogeneity-aware twitter bot detection with relational graph transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 3977–3985, 2022.
- [55] S. Feng, Z. Tan, H. Wan, N. Wang, Z. Chen, B. Zhang, Q. Zheng, W. Zhang, Z. Lei, S. Yang, et al. Twibot-22: Towards graph-based twitter bot detection. *Advances in Neural Information Processing Systems*, 35:35254–35269, 2022.
- [56] S. Feng, V. Balachandran, Y. Bai, and Y. Tsvetkov. Factkb: Generalizable factuality evaluation using language models enhanced with factual knowledge. [arXiv preprint arXiv:2305.08281](#), 2023.
- [57] S. Feng, C. Y. Park, Y. Liu, and Y. Tsvetkov. From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair nlp models. [arXiv preprint arXiv:2305.08283](#), 2023.
- [58] E. Ferrara. Social bot detection in the age of chatgpt: Challenges and opportunities. *First Monday*, 2023.
- [59] S. E. Finch, E. S. Paek, and J. D. Choi. Leveraging large language models for automated dialogue analysis. [arXiv preprint arXiv:2309.06490](#), 2023.
- [60] M. Foosherian, H. Purwins, P. Rathnayake, T. Alam, R. Teimao, and K.-D. Thoben. Enhancing pipeline-based conversational agents with large language models. [arXiv preprint arXiv:2309.03748](#), 2023.
- [61] L. Friedman, S. Ahuja, D. Allen, T. Tan, H. Sidahmed, C. Long, J. Xie, G. Schubiner, A. Patel, H. Lara, et al. Leveraging large language models in conversational recommender systems. [arXiv preprint arXiv:2305.07961](#), 2023.
- [62] G. Fu, Q. Zhao, J. Li, D. Luo, C. Song, W. Zhai, S. Liu, F. Wang, Y. Wang, L. Cheng, et al. Enhancing psychological counseling with large language model: A multifaceted decision-support system for non-professionals. [arXiv preprint arXiv:2308.15192](#), 2023.
- [63] Y. Gao, T. Sheng, Y. Xiang, Y. Xiong, H. Wang, and J. Zhang. Chat-rec: Towards interactive and explainable llms-augmented recommender system. [arXiv preprint arXiv:2303.14524](#), 2023.
- [64] H. Ghanadian, I. Nejadgholi, and H. A. Osman. Chatgpt for suicide risk assessment on social media: Quantitative evaluation of model performance, potentials and limitations. [arXiv preprint arXiv:2306.09390](#), 2023.
- [65] A. Grover and J. Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 855–864, 2016.

- [66] H. W. Hanley and Z. Durumeric. Machine-made media: Monitoring the mobilization of machine-generated articles on misinformation and mainstream news websites. [arXiv preprint arXiv:2305.09820](#), 2023.
- [67] Z. S. Harris. Distributional structure. *Word*, 10(2-3):146–162, 1954.
- [68] M. M. Hassan, A. Knipper, and S. K. K. Santu. Chatgpt as your personal data scientist. [arXiv preprint arXiv:2305.13657](#), 2023.
- [69] X. He, K. Deng, X. Wang, Y. Li, Y. Zhang, and M. Wang. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 639–648, 2020.
- [70] X. He, X. Bresson, T. Laurent, and B. Hooi. Explanations as features: Llm-based features for text-attributed graphs. [arXiv preprint arXiv:2305.19523](#), 2023.
- [71] Z. He, W. Liu, W. Guo, J. Qin, Y. Zhang, Y. Hu, and R. Tang. A survey on user behavior modeling in recommender systems. [arXiv preprint arXiv:2302.11087](#), 2023.
- [72] Z. He, Z. Xie, R. Jha, H. Steck, D. Liang, Y. Feng, B. P. Majumder, N. Kallus, and J. McAuley. Large language models as zero-shot conversational recommenders. [arXiv preprint arXiv:2308.10053](#), 2023.
- [73] Y. Hou, J. Zhang, Z. Lin, H. Lu, R. Xie, J. McAuley, and W. X. Zhao. Large language models are zero-shot rankers for recommender systems. [arXiv preprint arXiv:2305.08845](#), 2023.
- [74] E. J. Hu, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2021.
- [75] J. Hu, F. Jin, G. Zhang, J. Wang, and Y. Yang. A user profile modeling method based on word2vec. In *2017 IEEE International Conference on Software Quality, Reliability and Security Companion (QRS-C)*, pages 410–414. IEEE, 2017.
- [76] K. Hu, M. Yan, J. T. Zhou, I. W. Tsang, W. H. Chong, and Y. K. Yap. Ladder-of-thought: Using knowledge as steps to elevate stance detection. [arXiv preprint arXiv:2308.16763](#), 2023.
- [77] Z. Hu, Y. Dong, K. Wang, and Y. Sun. Heterogeneous graph transformer. In *Proceedings of the web conference 2020*, pages 2704–2710, 2020.
- [78] Z. Hu, Y. Feng, A. T. Luu, B. Hooi, and A. Lipani. Unlocking the potential of user feedback: Leveraging large language model as user simulator to enhance dialogue system. [arXiv preprint arXiv:2306.09821](#), 2023.
- [79] J. Huang, H. Shao, and K. C.-C. Chang. Are large pre-trained language models leaking your personal information? In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2038–2047, 2022.
- [80] X. Huang, J. Lian, Y. Lei, J. Yao, D. Lian, and X. Xie. Recommender ai agent: Integrating large language models for interactive recommendations. [arXiv preprint arXiv:2308.16505](#), 2023.
- [81] X. Huang, W. Ruan, W. Huang, G. Jin, Y. Dong, C. Wu, S. Bensalem, R. Mu, Y. Qi, X. Zhao, et al. A survey of safety and trustworthiness of large language models through the lens of verification and validation. [arXiv preprint arXiv:2305.11391](#), 2023.
- [82] Y. Huang and L. Sun. Harnessing the power of chatgpt in fake news: An in-depth exploration in generation, detection and explanation. [arXiv preprint arXiv:2310.05046](#), 2023.
- [83] V. Hudeček and O. Dušek. Are llms all you need for task-oriented dialogue? [arXiv preprint arXiv:2304.06556](#), 2023.
- [84] J. Huynh, C. Jiao, P. Gupta, S. Mehri, P. Bajaj, V. Chaudhary, and M. Eskenazi. Understanding the effectiveness of very large language models on dialog evaluation. [arXiv preprint arXiv:2301.12004](#), 2023.
- [85] J. Jang, S. Kim, B. Y. Lin, Y. Wang, J. Hessel, L. Zettlemoyer, H. Hajishirzi, Y. Choi, and P. Ammanabrolu. Personalized soups: Personalized large language model alignment via post-hoc parameter merging. [arXiv preprint arXiv:2310.11564](#), 2023.
- [86] J. Ji, Z. Li, S. Xu, W. Hua, Y. Ge, J. Tan, and Y. Zhang. Genrec: Large language model for generative recommendation. [arXiv e-prints](#), pages arXiv–2307, 2023.
- [87] Y. Ji, W. Wu, H. Zheng, Y. Hu, X. Chen, and L. He. Is chatgpt a good personality recognizer? a preliminary study. [arXiv preprint arXiv:2307.03952](#), 2023.
- [88] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. J. Bang, A. Madotto, and P. Fung. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38, 2023.
- [89] B. Jiang, Z. Tan, A. Nirmal, and H. Liu. Disinformation detection: An evolving challenge in the age of llms. [arXiv](#)

preprint arXiv:2309.15847, 2023.

- [90] H. Jiang, X. Zhang, X. Cao, J. Kabbara, and D. Roy. Personllm: Investigating the ability of gpt-3.5 to express personality traits and gender differences. [arXiv preprint arXiv:2305.02547](#), 2023.
- [91] M. Jiang, P. Cui, R. Liu, Q. Yang, F. Wang, W. Zhu, and S. Yang. Social contextual recommendation. In *Proceedings of the 21st ACM international conference on Information and knowledge management*, pages 45–54, 2012.
- [92] M. Jiang, P. Cui, A. Beutel, C. Faloutsos, and S. Yang. Catchsync: catching synchronized behavior in large directed graphs. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 941–950, 2014.
- [93] M. Jiang, P. Cui, A. Beutel, C. Faloutsos, and S. Yang. Detecting suspicious following behavior in multimillion-node social networks. In *Proceedings of the 23rd International Conference on World Wide Web*, pages 305–306, 2014.
- [94] M. Jiang, P. Cui, F. Wang, W. Zhu, and S. Yang. Scalable recommendation with social contextual information. *IEEE Transactions on Knowledge and Data Engineering*, 26(11):2789–2802, 2014.
- [95] M. Jiang, P. Cui, and C. Faloutsos. Suspicious behavior detection: Current trends and future directions. *IEEE intelligent systems*, 31(1):31–39, 2016.
- [96] M. Jiang, C. Faloutsos, and J. Han. Catchtartan: Representing and summarizing dynamic multicontextual behaviors. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 945–954, 2016.
- [97] Y. Jiang, R. Qiu, Y. Zhang, and P.-F. Zhang. Balanced and explainable social media analysis for public health with large language models. [arXiv preprint arXiv:2309.05951](#), 2023.
- [98] I. Joshi, R. Budhiraja, P. D. Tanna, L. Jain, M. Deshpande, A. Srivastava, S. Rallapalli, H. D. Akolekar, J. S. Challa, and D. Kumar. From "let's google" to "let's chatgpt": Student and instructor perspectives on the influence of llms on undergraduate engineering education. [arXiv preprint arXiv:2309.10694](#), 2023.
- [99] J. Kaddour, J. Harris, M. Mozes, H. Bradley, R. Raileanu, and R. McHardy. Challenges and applications of large language models. [arXiv preprint arXiv:2307.10169](#), 2023.
- [100] W.-C. Kang, J. Ni, N. Mehta, M. Sathiamoorthy, L. Hong, E. Chi, and D. Z. Cheng. Do llms understand user preferences? evaluating llms on user rating prediction. [arXiv preprint arXiv:2305.06474](#), 2023.
- [101] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei. Scaling laws for neural language models. [arXiv preprint arXiv:2001.08361](#), 2020.
- [102] J. D. M.-W. C. Kenton and L. K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019.
- [103] M. Khalil and E. Er. Will chatgpt get you caught? rethinking of plagiarism detection. [arXiv preprint arXiv:2302.04335](#), 2023.
- [104] K. Kheiri and H. Karimi. Sentimentgpt: Exploiting gpt for advanced sentiment analysis and its departure from current machine learning. [arXiv preprint arXiv:2307.10234](#), 2023.
- [105] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2016.
- [106] C. Kong, Y. Fan, X. Wan, F. Jiang, and B. Wang. Large language model as a user simulator. [arXiv preprint arXiv:2308.11534](#), 2023.
- [107] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8): 30–37, 2009.
- [108] M. Kosinski. Theory of mind may have spontaneously emerged in large language models. [arXiv preprint arXiv:2302.02083](#), 2023.
- [109] C. Koyuturk, M. Yavari, E. Theophilou, S. Bursic, G. Donabauer, A. Telari, A. Testa, R. Boiano, A. Gabbiadini, D. Hernandez-Leo, et al. Developing effective educational chatbots with chatgpt prompts: Insights from preliminary tests in a case study on social media literacy. [arXiv preprint arXiv:2306.10645](#), 2023.
- [110] M. Labonne and S. Moran. Spam-t5: Benchmarking large language models for few-shot email spam detection. [arXiv preprint arXiv:2304.01238](#), 2023.
- [111] T. Lai, Y. Shi, Z. Du, J. Wu, K. Fu, Y. Dou, and Z. Wang. Psy-llm: Scaling up global mental health psychological services with ai-based large language models. [arXiv preprint arXiv:2307.11991](#), 2023.

- [112] K. Lakkaraju, S. K. R. Vuruma, V. Pallagani, B. Muppasani, and B. Srivastava. Can llms be good financial advisors?: An initial study in personal decision making for optimized outcomes. [arXiv preprint arXiv:2307.07422](#), 2023.
- [113] D. Lee and H. S. Seung. Algorithms for non-negative matrix factorization. [Advances in neural information processing systems](#), 13, 2000.
- [114] J. A. Leite, O. Razuvayevskaya, K. Bontcheva, and C. Scarton. Detecting misinformation with llm-predicted credibility signals and weak supervision. [arXiv preprint arXiv:2309.07601](#), 2023.
- [115] B. Lester, R. Al-Rfou, and N. Constant. The power of scale for parameter-efficient prompt tuning. In [Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing](#), pages 3045–3059, 2021.
- [116] Y. Levine, N. Wies, D. Jannai, D. Navon, Y. Hoshen, and A. Shashua. The inductive bias of in-context learning: Rethinking pretraining example design. [arXiv preprint arXiv:2110.04541](#), 2021.
- [117] C. Li, M. Zhang, Q. Mei, Y. Wang, S. A. Hombaiyah, Y. Liang, and M. Bendersky. Teach llms to personalize—an approach inspired by writing education. [arXiv preprint arXiv:2308.07968](#), 2023.
- [118] H. Li, Y. Wu, V. Schlegel, R. Batista-Navarro, T.-T. Nguyen, A. R. Kashyap, X. Zeng, D. Beck, S. Winkler, and G. Nenadic. Pulsar: Pre-training with extracted healthcare terms for summarising patients’ problems and data augmentation with black-box large language models. [arXiv preprint arXiv:2306.02754](#), 2023.
- [119] J. Li, T. Tang, W. X. Zhao, and J.-R. Wen. Pretrained language models for text generation: A survey.
- [120] J. Li, W. Zhang, T. Wang, G. Xiong, A. Lu, and G. Medioni. Gpt4rec: A generative framework for personalized recommendation and user interests interpretation. [arXiv preprint arXiv:2304.03879](#), 2023.
- [121] N. Li, B. Kang, and T. De Bie. Llm4jobs: Unsupervised occupation extraction and standardization leveraging large language models. [arXiv preprint arXiv:2309.09708](#), 2023.
- [122] P. Li, Y. Wang, E. H. Chi, and M. Chen. Prompt tuning large language models on personalized aspect extraction for recommendations. [arXiv preprint arXiv:2306.01475](#), 2023.
- [123] R. Li, W. Deng, Y. Cheng, Z. Yuan, J. Zhang, and F. Yuan. Exploring the upper limits of text-based collaborative filtering using large language models: Discoveries and insights. [arXiv preprint arXiv:2305.11700](#), 2023.
- [124] S. Li and H. Zhao. A survey on representation learning for user modeling. In [Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence](#), pages 4997–5003, 2021.
- [125] X. Li, Y. Zhang, and E. C. Malthouse. A preliminary study of chatgpt on news recommendation: Personalization, provider fairness, fake news. [arXiv preprint arXiv:2306.10702](#), 2023.
- [126] X. L. Li and P. Liang. Prefix-tuning: Optimizing continuous prompts for generation. [arXiv preprint arXiv:2101.00190](#), 2021.
- [127] Y. Li, Z. Li, K. Zhang, R. Dan, S. Jiang, and Y. Zhang. Chatdoctor: A medical chat model fine-tuned on a large language model meta-ai (llama) using medical domain knowledge. [Cureus](#), 15(6), 2023.
- [128] Z. Li, B. Peng, P. He, and X. Yan. Do you really follow me? adversarial instructions for evaluating the robustness of large language models. [arXiv preprint arXiv:2308.10819](#), 2023.
- [129] G. Lin and Y. Zhang. Sparks of artificial general recommender (agr): Early experiments with chatgpt. [arXiv preprint arXiv:2305.04518](#), 2023.
- [130] J. Lin, X. Dai, Y. Xi, W. Liu, B. Chen, X. Li, C. Zhu, H. Guo, Y. Yu, R. Tang, et al. How can recommender systems benefit from large language models: A survey. [arXiv preprint arXiv:2306.05817](#), 2023.
- [131] J. Lin, R. Shan, C. Zhu, K. Du, B. Chen, S. Quan, R. Tang, Y. Yu, and W. Zhang. Rella: Retrieval-enhanced large language models for lifelong sequential behavior comprehension in recommendation. [arXiv preprint arXiv:2308.11131](#), 2023.
- [132] Y.-T. Lin and Y.-N. Chen. Llm-eval: Unified multi-dimensional automatic evaluation for open-domain conversations with large language models. [arXiv preprint arXiv:2305.13711](#), 2023.
- [133] J. Liu, C. Liu, R. Lv, K. Zhou, and Y. Zhang. Is chatgpt a good recommender? a preliminary study. [arXiv preprint arXiv:2304.10149](#), 2023.
- [134] J. Liu, C. Liu, P. Zhou, Q. Ye, D. Chong, K. Zhou, Y. Xie, Y. Cao, S. Wang, C. You, et al. Llmrec: Benchmarking large language models on recommendation task. [arXiv preprint arXiv:2308.12241](#), 2023.
- [135] Q. Liu, N. Chen, T. Sakai, and X.-M. Wu. A first look at llm-powered generative news recommendation. [arXiv preprint arXiv:2305.06566](#), 2023.

- [136] Q. Liu, N. Chen, T. Sakai, and X.-M. Wu. Once: Boosting content-based recommendation with both open- and closed-source large language models, 2023.
- [137] T. Liu, Y. Zhang, C. Brockett, Y. Mao, Z. Sui, W. Chen, and B. Dolan. A token-level reference-free hallucination detection benchmark for free-form text generation. [arXiv preprint arXiv:2104.08704](#), 2021.
- [138] X. Liu, Z. Zhang, Y. Wang, Y. Lan, and C. Shen. Coco: Coherence-enhanced machine-generated text detection under data limitation with contrastive learning. [arXiv preprint arXiv:2212.10341](#), 2022.
- [139] X. Liu, D. McDuff, G. Kovacs, I. Galatzer-Levy, J. Sunshine, J. Zhan, M.-Z. Poh, S. Liao, P. Di Achille, and S. Patel. Large language models are few-shot health learners. [arXiv preprint arXiv:2305.15525](#), 2023.
- [140] X. Liu, H. Yu, H. Zhang, Y. Xu, X. Lei, H. Lai, Y. Gu, H. Ding, K. Men, K. Yang, et al. Agentbench: Evaluating llms as agents. [arXiv preprint arXiv:2308.03688](#), 2023.
- [141] X. Liu, Y. Zheng, Z. Du, M. Ding, Y. Qian, Z. Yang, and J. Tang. Gpt understands, too. [AI Open](#), 2023.
- [142] X.-Y. Liu, G. Wang, and D. Zha. Fingpt: Democratizing internet-scale data for financial large language models. [arXiv preprint arXiv:2307.10485](#), 2023.
- [143] Z. Liu, Z. Wu, M. Hu, B. Zhao, L. Zhao, T. Zhang, H. Dai, X. Chen, Y. Shen, S. Li, et al. Pharmacygpt: The ai pharmacist. [arXiv preprint arXiv:2307.10432](#), 2023.
- [144] N. Lu, S. Liu, R. He, and K. Tang. Large language models can be guided to evade ai-generated text detection. [arXiv preprint arXiv:2305.10847](#), 2023.
- [145] H. Lyu, S. Jiang, H. Zeng, Y. Xia, and J. Luo. Llm-rec: Personalized recommendation via prompting large language models. [arXiv preprint arXiv:2307.15780](#), 2023.
- [146] J. Manyika. An overview of bard: an early experiment with generative ai. Technical report, Tech. rep., Technical report, Google AI, 2023.
- [147] S. Matwin, A. Milios, P. Prałat, A. Soares, and F. Théberge. Survey of generative methods for social media analysis. [arXiv preprint arXiv:2112.07041](#), 2021.
- [148] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. [arXiv preprint arXiv:1301.3781](#), 2013.
- [149] S. Min, K. Krishna, X. Lyu, M. Lewis, W.-t. Yih, P. W. Koh, M. Iyyer, L. Zettlemoyer, and H. Hajishirzi. Factscore: Fine-grained atomic evaluation of factual precision in long form text generation. [arXiv preprint arXiv:2305.14251](#), 2023.
- [150] L. Mo, S. Chen, Z. Chen, X. Deng, A. Lewis, S. Singh, S. Stevens, C.-Y. Tai, Z. Wang, X. Yue, et al. Roll up your sleeves: Working with a collaborative and engaging task-oriented dialogue system. [arXiv preprint arXiv:2307.16081](#), 2023.
- [151] J. X. Morris, V. Kuleshov, V. Shmatikov, and A. M. Rush. Text embeddings reveal (almost) as much as text. [arXiv preprint arXiv:2310.06816](#), 2023.
- [152] M. Mozes, X. He, B. Kleinberg, and L. D. Griffin. Use of llms for illicit purposes: Threats, prevention measures, and vulnerabilities. [arXiv preprint arXiv:2308.12833](#), 2023.
- [153] Y. Mu, B. P. Wu, W. Thorne, A. Robinson, N. Aletras, C. Scarton, K. Bontcheva, and X. Song. Navigating prompt complexity for zero-shot classification: A study of large language models in computational social science. [arXiv preprint arXiv:2305.14310](#), 2023.
- [154] S. Mysore, A. McCallum, and H. Zamani. Large language model augmented narrative driven recommendations. [arXiv preprint arXiv:2306.02250](#), 2023.
- [155] S. Natkin and C. Yan. User model in multiplayer mixed reality entertainment applications. In [Proceedings of the 2006 ACM SIGCHI international conference on Advances in computer entertainment technology](#), pages 74–es, 2006.
- [156] T. T. Nguyen, C. Wilson, and J. Dalins. Fine-tuning llama 2 large language models for detecting online sexual predatory chats and abusive texts. [arXiv preprint arXiv:2308.14683](#), 2023.
- [157] P. Ochieng. Are large language models fit for guided reading? [arXiv preprint arXiv:2305.10645](#), 2023.
- [158] A. Olga, A. Saini, G. Zapata, D. Searsmith, B. Cope, M. Kalantzis, V. Castro, T. Kourkoulou, J. Jones, R. A. da Silva, et al. Generative ai: Implications and applications for education. [arXiv preprint arXiv:2305.07605](#), 2023.
- [159] OpenAI. Gpt-4 technical report, 2023.
- [160] M. S. Orenstrakh, O. Karnalim, C. A. Suarez, and M. Liut. Detecting llm-generated text in computing education: A

- comparative study for chatgpt cases. [arXiv preprint arXiv:2307.07411](#), 2023.
- [161] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- [162] L. PAGE. The pagerank citation ranking: Bringing order to the web. In *Proc. of the 7th WWW Conf.*, 1998, 1998.
- [163] S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, and X. Wu. Unifying large language models and knowledge graphs: A roadmap. [arXiv preprint arXiv:2306.08302](#), 2023.
- [164] Y. Pan, L. Pan, W. Chen, P. Nakov, M.-Y. Kan, and W. Y. Wang. On the risk of misinformation pollution with large language models. [arXiv preprint arXiv:2305.13661](#), 2023.
- [165] S. Parikh, Q. Vohra, P. Tumbade, and M. Tiwari. Exploring zero and few-shot techniques for intent classification. [arXiv preprint arXiv:2305.07157](#), 2023.
- [166] J. S. Park, J. C. O’Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein. Generative agents: Interactive simulacra of human behavior. [arXiv preprint arXiv:2304.03442](#), 2023.
- [167] B. M. Pavlyshenko. Analysis of disinformation and fake news detection using fine-tuned large language model. [arXiv preprint arXiv:2309.04704](#), 2023.
- [168] A. Pegoraro, K. Kumari, H. Fereidooni, and A.-R. Sadeghi. To chatgpt, or not to chatgpt: That is the question! [arXiv preprint arXiv:2304.01487](#), 2023.
- [169] B. Peng, E. Alcaide, Q. Anthony, A. Albalak, S. Arcadinho, H. Cao, X. Cheng, M. Chung, M. Grella, K. K. GV, et al. Rwkv: Reinventing rnns for the transformer era. [arXiv preprint arXiv:2305.13048](#), 2023.
- [170] L. Peng, Y. Zhang, and J. Shang. Generating efficient training data via llm-based attribute manipulation. [arXiv preprint arXiv:2307.07099](#), 2023.
- [171] W. Peng, D. Xu, T. Xu, J. Zhang, and E. Chen. Are gpt embeddings useful for ads and recommendation? In *International Conference on Knowledge Science, Engineering and Management*, pages 151–162. Springer, 2023.
- [172] B. Perozzi, R. Al-Rfou, and S. Skiena. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 701–710, 2014.
- [173] H. Peters and S. Matz. Large language models can infer psychological dispositions of social media users. [arXiv preprint arXiv:2309.08631](#), 2023.
- [174] T. Phung, V.-A. Pădurean, J. Cambronero, S. Gulwani, T. Kohn, R. Majumdar, A. Singla, and G. Soares. Generative ai for programming education: Benchmarking chatgpt, gpt-4, and human tutors. *International Journal of Management*, 21(2):100790, 2023.
- [175] R. Pope, S. Douglas, A. Chowdhery, J. Devlin, J. Bradbury, J. Heek, K. Xiao, S. Agrawal, and J. Dean. Efficiently scaling transformer inference. *Proceedings of Machine Learning and Systems*, 5, 2023.
- [176] X. Pu, M. Gao, and X. Wan. Summarization is (almost) dead. [arXiv preprint arXiv:2309.09558](#), 2023.
- [177] H. Qi, Q. Zhao, C. Song, W. Zhai, D. Luo, S. Liu, Y. J. Yu, F. Wang, H. Zou, B. X. Yang, et al. Evaluating the efficacy of supervised learning vs large language models for identifying cognitive distortions and suicidal risks in chinese social media. [arXiv preprint arXiv:2309.03564](#), 2023.
- [178] W. Qin, Z. Chen, L. Wang, Y. Lan, W. Ren, and R. Hong. Read, diagnose and chat: Towards explainable and interactive llms-augmented depression detection in social media. [arXiv preprint arXiv:2305.05138](#), 2023.
- [179] X. Qiu, T. Sun, Y. Xu, Y. Shao, N. Dai, and X. Huang. Pre-trained models for natural language processing: A survey. *Science China Technological Sciences*, 63(10):1872–1897, 2020.
- [180] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever, et al. Improving language understanding by generative pre-training.
- [181] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [182] 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*, 21(1): 5485–5551, 2020.
- [183] A. Rao, S. Vashistha, A. Naik, S. Aditya, and M. Choudhury. Tricking llms into disobedience: Understanding,

- analyzing, and preventing jailbreaks. [arXiv preprint arXiv:2305.14965](#), 2023.
- [184] H. Rao, C. Leung, and C. Miao. Can chatgpt assess human personalities? a general evaluation framework. [arXiv preprint arXiv:2303.01248](#), 2023.
- [185] C. Richardson, Y. Zhang, K. Gillespie, S. Kar, A. Singh, Z. Raeesy, O. Z. Khan, and A. Sethy. Integrating summarization and retrieval for enhanced personalization via large language models. [arXiv preprint arXiv:2310.20081](#), 2023.
- [186] J. Robinson, C. M. Rytting, and D. Wingate. Leveraging large language models for multiple choice question answering. [arXiv preprint arXiv:2210.12353](#), 2022.
- [187] X. Runfeng, C. Xiangyang, Y. Zhou, W. Xin, X. Zhanwei, Z. Kai, et al. Lkpnr: Llm and kg for personalized news recommendation framework. [arXiv preprint arXiv:2308.12028](#), 2023.
- [188] S. Russell. Human-compatible artificial intelligence. [Human-like machine intelligence](#), pages 3–23, 2021.
- [189] A. Salemi, S. Mysore, M. Bendersky, and H. Zamani. Lamp: When large language models meet personalization. [arXiv preprint arXiv:2304.11406](#), 2023.
- [190] S. Sanner, K. Balog, F. Radlinski, B. Wedin, and L. Dixon. Large language models are competitive near cold-start recommenders for language-and item-based preferences. In [Proceedings of the 17th ACM Conference on Recommender Systems](#), pages 890–896, 2023.
- [191] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl. Analysis of recommendation algorithms for e-commerce. In [Proceedings of the 2nd ACM Conference on Electronic Commerce](#), pages 158–167, 2000.
- [192] T. L. Scao, A. Fan, C. Akiki, E. Pavlick, S. Ilić, D. Hesslow, R. Castagné, A. S. Luccioni, F. Yvon, M. Gallé, et al. Bloom: A 176b-parameter open-access multilingual language model. [arXiv preprint arXiv:2211.05100](#), 2022.
- [193] J. B. Schafer, D. Frankowski, J. Herlocker, and S. Sen. Collaborative filtering recommender systems. In [The adaptive web: methods and strategies of web personalization](#), pages 291–324. Springer.
- [194] J. B. Schafer, J. A. Konstan, and J. Riedl. E-commerce recommendation applications. [Data mining and knowledge discovery](#), 5:115–153, 2001.
- [195] T. Schick, J. Dwivedi-Yu, R. Dessì, R. Raileanu, M. Lomeli, L. Zettlemoyer, N. Cancedda, and T. Scialom. Toolformer: Language models can teach themselves to use tools. [arXiv preprint arXiv:2302.04761](#), 2023.
- [196] V. Schlegel, H. Li, Y. Wu, A. Subramanian, T.-T. Nguyen, A. R. Kashyap, D. Beck, X. Zeng, R. T. Batista-Navarro, S. Winkler, et al. Pulsar at mediqa-sum 2023: Large language models augmented by synthetic dialogue convert patient dialogues to medical records. [arXiv preprint arXiv:2307.02006](#), 2023.
- [197] P. Sharma, K. Thapa, P. Dhakal, M. D. Upadhaya, S. Adhikari, and S. R. Khanal. Performance of chatgpt on usmle: Unlocking the potential of large language models for ai-assisted medical education. [arXiv preprint arXiv:2307.00112](#), 2023.
- [198] Y. Shen, K. Song, X. Tan, D. Li, W. Lu, and Y. Zhuang. Hugginggpt: Solving ai tasks with chatgpt and its friends in huggingface. [arXiv preprint arXiv:2303.17580](#), 2023.
- [199] C. Shi, Y. Li, J. Zhang, Y. Sun, and S. Y. Philip. A survey of heterogeneous information network analysis. [IEEE Transactions on Knowledge and Data Engineering](#), 29(1):17–37, 2016.
- [200] A. D. Shukla, L. Agarwal, J. Mein, R. Agarwal, et al. Catch me if you can: Identifying fraudulent physician reviews with large language models using generative pre-trained transformers. [arXiv preprint arXiv:2304.09948](#), 2023.
- [201] C. Si, D. Friedman, N. Joshi, S. Feng, D. Chen, and H. He. Measuring inductive biases of in-context learning with underspecified demonstrations. [arXiv preprint arXiv:2305.13299](#), 2023.
- [202] A. P. Singh and G. J. Gordon. Relational learning via collective matrix factorization. In [Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining](#), pages 650–658, 2008.
- [203] J. Su, T. Y. Zhuo, J. Mansurov, D. Wang, and P. Nakov. Fake news detectors are biased against texts generated by large language models. [arXiv preprint arXiv:2309.08674](#), 2023.
- [204] F. Sun, J. Liu, J. Wu, C. Pei, X. Lin, W. Ou, and P. Jiang. Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer. In [Proceedings of the 28th ACM international conference on information and knowledge management](#), pages 1441–1450, 2019.
- [205] K. Sun, Y. E. Xu, H. Zha, Y. Liu, and X. L. Dong. Head-to-tail: How knowledgeable are large language models (llm)? aka will llms replace knowledge graphs? [arXiv preprint arXiv:2308.10168](#), 2023.

- [206] X. Sun, X. Li, J. Li, F. Wu, S. Guo, T. Zhang, and G. Wang. Text classification via large language models. [arXiv preprint arXiv:2305.08377](#), 2023.
- [207] E. Svikhnushina and P. Pu. Approximating human evaluation of social chatbots with prompting. [arXiv preprint arXiv:2304.05253](#), 2023.
- [208] Z. Tan, S. Feng, M. Sclar, H. Wan, M. Luo, Y. Choi, and Y. Tsvetkov. Botpercent: Estimating twitter bot populations from groups to crowds. [arXiv preprint arXiv:2302.00381](#), 2023.
- [209] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei. Line: Large-scale information network embedding. In [Proceedings of the 24th international conference on world wide web](#), pages 1067–1077, 2015.
- [210] R. Tang, Y.-N. Chuang, and X. Hu. The science of detecting llm-generated texts. [arXiv preprint arXiv:2303.07205](#), 2023.
- [211] R. Tang, X. Han, X. Jiang, and X. Hu. Does synthetic data generation of llms help clinical text mining? [arXiv preprint arXiv:2303.04360](#), 2023.
- [212] K. Tao, F. Abel, Q. Gao, and G.-J. Houben. Tums: twitter-based user modeling service. In [The Semantic Web: ESWC 2011 Workshops: ESWC 2011 Workshops, Heraklion, Greece, May 29-30, 2011, Revised Selected Papers 8](#), pages 269–283. Springer, 2012.
- [213] R. Taori, I. Gulrajani, T. Zhang, Y. Dubois, X. Li, C. Guestrin, P. Liang, and T. B. Hashimoto. Alpaca: A strong, replicable instruction-following model. [Stanford Center for Research on Foundation Models. https://crfm.stanford.edu/2023/03/13/alpaca.html](#), 3(6):7, 2023.
- [214] H. Tian, W. Lu, T. O. Li, X. Tang, S.-C. Cheung, J. Klein, and T. F. Bissyandé. Is chatgpt the ultimate programming assistant—how far is it? [arXiv preprint arXiv:2304.11938](#), 2023.
- [215] X. Tie, M. Shin, A. Pirasteh, N. Ibrahim, Z. Huemann, S. M. Castellino, K. M. Kelly, J. Garrett, J. Hu, S. Y. Cho, et al. Automatic personalized impression generation for pet reports using large language models. [arXiv preprint arXiv:2309.10066](#), 2023.
- [216] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, et al. Llama: Open and efficient foundation language models. [arXiv preprint arXiv:2302.13971](#), 2023.
- [217] Q. Tu, C. Chen, J. Li, Y. Li, S. Shang, D. Zhao, R. Wang, and R. Yan. Characterchat: Learning towards conversational ai with personalized social support. [arXiv preprint arXiv:2308.10278](#), 2023.
- [218] X. Tu, J. Zou, W. J. Su, and L. Zhang. What should data science education do with large language models? [arXiv preprint arXiv:2307.02792](#), 2023.
- [219] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. [Advances in neural information processing systems](#), 30, 2017.
- [220] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. In [International Conference on Learning Representations](#), 2018.
- [221] V. Veselovsky, M. H. Ribeiro, A. Arora, M. Josifoski, A. Anderson, and R. West. Generating faithful synthetic data with large language models: A case study in computational social science. [arXiv preprint arXiv:2305.15041](#), 2023.
- [222] V. Viswanathan, C. Zhao, A. Bertsch, T. Wu, and G. Neubig. Prompt2model: Generating deployable models from natural language instructions. [arXiv preprint arXiv:2308.12261](#), 2023.
- [223] J. P. Wahle, T. Ruas, F. Kirstein, and B. Gipp. How large language models are transforming machine-paraphrased plagiarism. [arXiv preprint arXiv:2210.03568](#), 2022.
- [224] D. Wang, M. Jiang, M. Syed, O. Conway, V. Juneja, S. Subramanian, and N. V. Chawla. Calendar graph neural networks for modeling time structures in spatiotemporal user behaviors. In [Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining](#), pages 2581–2589, 2020.
- [225] D. Wang, Z. Zhang, Y. Ma, T. Zhao, T. Jiang, N. Chawla, and M. Jiang. Modeling co-evolution of attributed and structural information in graph sequence. [IEEE Transactions on Knowledge and Data Engineering](#), 2021.
- [226] H. Wang, S. Feng, T. He, Z. Tan, X. Han, and Y. Tsvetkov. Can language models solve graph problems in natural language? [arXiv preprint arXiv:2305.10037](#), 2023.
- [227] H. Wang, M. S. Hee, M. R. Awal, K. T. W. Choo, and R. K.-W. Lee. Evaluating gpt-3 generated explanations for hateful content moderation. [arXiv preprint arXiv:2305.17680](#), 2023.
- [228] J. Wang, Z. Yao, A. Mitra, S. Osebe, Z. Yang, and H. Yu. Umass_bionlp at mediqa-chat 2023: Can llms generate

- high-quality synthetic note-oriented doctor-patient conversations? [arXiv preprint arXiv:2306.16931](#), 2023.
- [229] L. Wang and E.-P. Lim. Zero-shot next-item recommendation using large pretrained language models. [arXiv preprint arXiv:2304.03153](#), 2023.
- [230] L. Wang, C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. Chen, Y. Lin, et al. A survey on large language model based autonomous agents. [arXiv preprint arXiv:2308.11432](#), 2023.
- [231] L. Wang, J. Zhang, X. Chen, Y. Lin, R. Song, W. X. Zhao, and J.-R. Wen. Recagent: A novel simulation paradigm for recommender systems. [arXiv preprint arXiv:2306.02552](#), 2023.
- [232] T. Wang, A. Roberts, D. Hesslow, T. Le Scao, H. W. Chung, I. Beltagy, J. Launay, and C. 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.
- [233] W. Wang, X. Lin, F. Feng, X. He, and T.-S. Chua. Generative recommendation: Towards next-generation recommender paradigm. [arXiv preprint arXiv:2304.03516](#), 2023.
- [234] X. Wang, J. Wei, D. Schuurmans, Q. Le, E. Chi, S. Narang, A. Chowdhery, and D. Zhou. Self-consistency improves chain of thought reasoning in language models. [arXiv preprint arXiv:2203.11171](#), 2022.
- [235] X. Wang, X. Tang, W. X. Zhao, J. Wang, and J.-R. Wen. Rethinking the evaluation for conversational recommendation in the era of large language models. [arXiv preprint arXiv:2305.13112](#), 2023.
- [236] Y. Wang, Z. Jiang, Z. Chen, F. Yang, Y. Zhou, E. Cho, X. Fan, X. Huang, Y. Lu, and Y. Yang. Recmind: Large language model powered agent for recommendation. [arXiv preprint arXiv:2308.14296](#), 2023.
- [237] Y. Wang, Y. Zhao, and L. Petzold. Are large language models ready for healthcare? a comparative study on clinical language understanding. [arXiv preprint arXiv:2304.05368](#), 2023.
- [238] Y. Wang, W. Zhong, L. Li, F. Mi, X. Zeng, W. Huang, L. Shang, X. Jiang, and Q. Liu. Aligning large language models with human: A survey. [arXiv preprint arXiv:2307.12966](#), 2023.
- [239] J. Wei, M. Bosma, V. Y. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, and Q. V. Le. Finetuned language models are zero-shot learners. [arXiv preprint arXiv:2109.01652](#), 2021.
- [240] J. Wei, Y. Tay, R. Bommasani, C. Raffel, B. Zoph, S. Borgeaud, D. Yogatama, M. Bosma, D. Zhou, D. Metzler, et al. Emergent abilities of large language models. [arXiv preprint arXiv:2206.07682](#), 2022.
- [241] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. [Advances in Neural Information Processing Systems](#), 35:24824–24837, 2022.
- [242] L. Wu, Z. Qiu, Z. Zheng, H. Zhu, and E. Chen. Exploring large language model for graph data understanding in online job recommendations. [arXiv preprint arXiv:2307.05722](#), 2023.
- [243] L. Wu, Z. Zheng, Z. Qiu, H. Wang, H. Gu, T. Shen, C. Qin, C. Zhu, H. Zhu, Q. Liu, et al. A survey on large language models for recommendation. [arXiv preprint arXiv:2305.19860](#), 2023.
- [244] P. Y. Wu, J. A. Tucker, J. Nagler, and S. Messing. Large language models can be used to estimate the ideologies of politicians in a zero-shot learning setting. [arXiv preprint arXiv:2303.12057](#), 2023.
- [245] S. Wu, O. Irsoy, S. Lu, V. Dabrovolski, M. Dredze, S. Gehrmann, P. Kambadur, D. Rosenberg, and G. Mann. Bloomberggpt: A large language model for finance. [arXiv preprint arXiv:2303.17564](#), 2023.
- [246] Y. Xi, W. Liu, J. Lin, J. Zhu, B. Chen, R. Tang, W. Zhang, R. Zhang, and Y. Yu. Towards open-world recommendation with knowledge augmentation from large language models. [arXiv preprint arXiv:2306.10933](#), 2023.
- [247] Z. Xi, W. Chen, X. Guo, W. He, Y. Ding, B. Hong, M. Zhang, J. Wang, S. Jin, E. Zhou, et al. The rise and potential of large language model based agents: A survey. [arXiv preprint arXiv:2309.07864](#), 2023.
- [248] X. Xu, B. Yao, Y. Dong, S. Gabriel, H. Yu, J. Hendler, M. Ghassemi, A. K. Dey, and D. Wang. Mental-llm: Leveraging large language models for mental health prediction via online text data. [arXiv preprint arXiv:2307.14385](#), 2023.
- [249] L. Yan, L. Sha, L. Zhao, Y. Li, R. Martinez-Maldonado, G. Chen, X. Li, Y. Jin, and D. Gašević. Practical and ethical challenges of large language models in education: A systematic literature review. [arXiv preprint arXiv:2303.13379](#), 2023.
- [250] Q. Yan, Y. Zhang, Q. Liu, S. Wu, and L. Wang. Relation-aware heterogeneous graph for user profiling. In [Proceedings of the 30th ACM International Conference on Information & Knowledge Management](#), pages 3573–3577, 2021.
- [251] C. Yang, X. Wang, Y. Lu, H. Liu, Q. V. Le, D. Zhou, and X. Chen. Large language models as optimizers. [arXiv](#)

- preprint [arXiv:2309.03409](https://arxiv.org/abs/2309.03409), 2023.
- [252] D. Yang, R. Yuan, Y. Fan, Y. Yang, Z. Wang, S. Wang, and H. Zhao. Refgpt: Reference-> truthful & customized dialogues generation by gpts and for gpts. [arXiv preprint arXiv:2305.14994](https://arxiv.org/abs/2305.14994), 2023.
- [253] K.-C. Yang and F. Menczer. Anatomy of an ai-powered malicious social botnet. [arXiv preprint arXiv:2307.16336](https://arxiv.org/abs/2307.16336), 2023.
- [254] L. Yang, F. Jiang, and H. Li. Is chatgpt involved in texts? measure the polish ratio to detect chatgpt-generated text. [arXiv preprint arXiv:2307.11380](https://arxiv.org/abs/2307.11380), 2023.
- [255] S. Yang, H. Zhao, S. Zhu, G. Zhou, H. Xu, Y. Jia, and H. Zan. Zhongjing: Enhancing the chinese medical capabilities of large language model through expert feedback and real-world multi-turn dialogue. [arXiv preprint arXiv:2308.03549](https://arxiv.org/abs/2308.03549), 2023.
- [256] S. Yao, D. Yu, J. Zhao, I. Shafran, T. L. Griffiths, Y. Cao, and K. Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. [arXiv preprint arXiv:2305.10601](https://arxiv.org/abs/2305.10601), 2023.
- [257] Y. Yao, Z. Li, and H. Zhao. Beyond chain-of-thought, effective graph-of-thought reasoning in large language models. [arXiv preprint arXiv:2305.16582](https://arxiv.org/abs/2305.16582), 2023.
- [258] Y. Yao, P. Wang, B. Tian, S. Cheng, Z. Li, S. Deng, H. Chen, and N. Zhang. Editing large language models: Problems, methods, and opportunities. [arXiv preprint arXiv:2305.13172](https://arxiv.org/abs/2305.13172), 2023.
- [259] R. Ye, C. Zhang, R. Wang, S. Xu, and Y. Zhang. Natural language is all a graph needs. [arXiv preprint arXiv:2308.07134](https://arxiv.org/abs/2308.07134), 2023.
- [260] B. Yin, J. Xie, Y. Qin, Z. Ding, Z. Feng, X. Li, and W. Lin. Heterogeneous knowledge fusion: A novel approach for personalized recommendation via llm. [arXiv preprint arXiv:2308.03333](https://arxiv.org/abs/2308.03333), 2023.
- [261] R. Ying, R. He, K. Chen, P. Eksombatchai, W. L. Hamilton, and J. Leskovec. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 974–983, 2018.
- [262] M. Yu, Z. Zhang, W. Yu, and M. Jiang. Pre-training language models for comparative reasoning. [arXiv preprint arXiv:2305.14457](https://arxiv.org/abs/2305.14457), 2023.
- [263] W. Yu, M. Yu, T. Zhao, and M. Jiang. Identifying referential intention with heterogeneous contexts. In *Proceedings of The Web Conference 2020*, pages 962–972, 2020.
- [264] W. Yu, C. Zhu, Z. Li, Z. Hu, Q. Wang, H. Ji, and M. Jiang. A survey of knowledge-enhanced text generation. *ACM Computing Surveys*, 54(11s):1–38, 2022.
- [265] X. Yu, Z. Chen, Y. Ling, S. Dong, Z. Liu, and Y. Lu. Temporal data meets llm—explainable financial time series forecasting. [arXiv preprint arXiv:2306.11025](https://arxiv.org/abs/2306.11025), 2023.
- [266] X. Yu, Y. Qi, K. Chen, G. Chen, X. Yang, P. Zhu, W. Zhang, and N. Yu. Gpt paternity test: Gpt generated text detection with gpt genetic inheritance. [arXiv preprint arXiv:2305.12519](https://arxiv.org/abs/2305.12519), 2023.
- [267] Y. Yu, Y. Zhuang, J. Zhang, Y. Meng, A. Ratner, R. Krishna, J. Shen, and C. Zhang. Large language model as attributed training data generator: A tale of diversity and bias. [arXiv preprint arXiv:2306.15895](https://arxiv.org/abs/2306.15895), 2023.
- [268] J. Yuan, R. Tang, X. Jiang, and X. Hu. Llm for patient-trial matching: Privacy-aware data augmentation towards better performance and generalizability. [arXiv preprint arXiv:2303.16756](https://arxiv.org/abs/2303.16756), 2023.
- [269] S. Yue, W. Chen, S. Wang, B. Li, C. Shen, S. Liu, Y. Zhou, Y. Xiao, S. Yun, W. Lin, et al. Disc-lawllm: Fine-tuning large language models for intelligent legal services. [arXiv preprint arXiv:2309.11325](https://arxiv.org/abs/2309.11325), 2023.
- [270] A. Zeng, X. Liu, Z. Du, Z. Wang, H. Lai, M. Ding, Z. Yang, Y. Xu, W. Zheng, X. Xia, et al. Glm-130b: An open bilingual pre-trained model. In *The Eleventh International Conference on Learning Representations*, 2022.
- [271] B. Zhang, D. Ding, and L. Jing. How would stance detection techniques evolve after the launch of chatgpt? [arXiv preprint arXiv:2212.14548](https://arxiv.org/abs/2212.14548), 2022.
- [272] D. Zhang, J. Yin, X. Zhu, and C. Zhang. Network representation learning: A survey. *IEEE transactions on Big Data*, 6(1):3–28, 2018.
- [273] J. Zhang. Graph-toolformer: To empower llms with graph reasoning ability via prompt augmented by chatgpt. [arXiv preprint arXiv:2304.11116](https://arxiv.org/abs/2304.11116), 2023.
- [274] J. Zhang, K. Bao, Y. Zhang, W. Wang, F. Feng, and X. He. Is chatgpt fair for recommendation? evaluating fairness in large language model recommendation. [arXiv preprint arXiv:2305.07609](https://arxiv.org/abs/2305.07609), 2023.

- [275] J. Zhang, K. Sun, A. Jagadeesh, M. Ghahfarokhi, D. Gupta, A. Gupta, V. Gupta, and Y. Guo. The potential and pitfalls of using a large language model such as chatgpt or gpt-4 as a clinical assistant. [arXiv preprint arXiv:2307.08152](#), 2023.
- [276] J. Zhang, R. Xie, Y. Hou, W. X. Zhao, L. Lin, and J.-R. Wen. Recommendation as instruction following: A large language model empowered recommendation approach. [arXiv preprint arXiv:2305.07001](#), 2023.
- [277] R. Zhang, J. Han, A. Zhou, X. Hu, S. Yan, P. Lu, H. Li, P. Gao, and Y. Qiao. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. [arXiv preprint arXiv:2303.16199](#), 2023.
- [278] S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. Diab, X. Li, X. V. Lin, et al. Opt: Open pre-trained transformer language models. [arXiv preprint arXiv:2205.01068](#), 2022.
- [279] Y. Zhang, Y. Li, L. Cui, D. Cai, L. Liu, T. Fu, X. Huang, E. Zhao, Y. Zhang, Y. Chen, et al. Siren’s song in the ai ocean: A survey on hallucination in large language models. [arXiv preprint arXiv:2309.01219](#), 2023.
- [280] Q. Zhao, Y. Zhang, D. Friedman, and F. Tan. E-commerce recommendation with personalized promotion. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pages 219–226, 2015.
- [281] T. Zhao, B. Ni, W. Yu, Z. Guo, N. Shah, and M. Jiang. Action sequence augmentation for early graph-based anomaly detection. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 2668–2678, 2021.
- [282] W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong, et al. A survey of large language models. [arXiv preprint arXiv:2303.18223](#), 2023.
- [283] Z. Zhao, E. Wallace, S. Feng, D. Klein, and S. Singh. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*, pages 12697–12706. PMLR, 2021.
- [284] C. Zheng, S. Sabour, J. Wen, and M. Huang. Augesc: Large-scale data augmentation for emotional support conversation with pre-trained language models. [arXiv preprint arXiv:2202.13047](#), 2022.
- [285] L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. [arXiv preprint arXiv:2306.05685](#), 2023.
- [286] Z. Zheng, L. Liao, Y. Deng, and L. Nie. Building emotional support chatbots in the era of llms. [arXiv preprint arXiv:2308.11584](#), 2023.
- [287] Z. Zheng, Z. Qiu, X. Hu, L. Wu, H. Zhu, and H. Xiong. Generative job recommendations with large language model. [arXiv preprint arXiv:2307.02157](#), 2023.
- [288] A. Zhiyuli, Y. Chen, X. Zhang, and X. Liang. Bookgpt: A general framework for book recommendation empowered by large language model. [arXiv preprint arXiv:2305.15673](#), 2023.
- [289] C. Zhou, P. Liu, P. Xu, S. Iyer, J. Sun, Y. Mao, X. Ma, A. Efrat, P. Yu, L. Yu, et al. Lima: Less is more for alignment. [arXiv preprint arXiv:2305.11206](#), 2023.
- [290] D. Zhou, N. Schärli, L. Hou, J. Wei, N. Scales, X. Wang, D. Schuurmans, C. Cui, O. Bousquet, Q. Le, et al. Least-to-most prompting enables complex reasoning in large language models. [arXiv preprint arXiv:2205.10625](#), 2022.
- [291] Y. Zhu, H. Wang, Y. Wang, Y. Li, Y. Yuan, and J. Qiang. Clickbait detection via large language models. [arXiv preprint arXiv:2306.09597](#), 2023.
- [292] C. Ziems, W. Held, O. Shaikh, J. Chen, Z. Zhang, and D. Yang. Can large language models transform computational social science? [arXiv preprint arXiv:2305.03514](#), 2023.
- [293] N. Ziems, G. Liu, J. Flanagan, and M. Jiang. Explaining tree model decisions in natural language for network intrusion detection. [arXiv preprint arXiv:2310.19658](#), 2023.

Table 5: Representative approaches using the LLMs-as-Predictors paradigm.

Roles	Applications	LLM Backbones	Models
Common Generative Reasoner	Recommendation	GPT family	Liu et al. [133], NIR [229], BookGPT [288], Hou et al. [73], Li et al. [125], Di Palma et al. [44]
		PaLM	Sanner et al. [190]
		Llama/Vicuna	PALR [26], GIRL [287], TallRec [8], GLRec [242], ReLLa [131]
		FLAN-T5	Zhang et al. [276]
	Intelligent Assistant	GPT family	LaMP [189], Chakrabarty et al. [20]
		FLAN-T5/T5	LaMP [189], Li et al. [117]
		ChatGLM	Disc-LawLLM [269]
	User Profiling	ChatGPT	Rao et al. [184], Ji et al. [87], Wu et al. [244]
		LaMDA, PaLM	Christakopoulou et al. [32]
	Dialogue System	ChatGPT	Fan and Jiang [51]
	Education	ChatGPT	Sharma et al. [197], Elkins et al. [48], Ochieng [157], C-LLM [158], Phung et al. [174]
		BARD	Ochieng [157]
	Healthcare	GPT family	Wang et al. [237], Ghanadian et al. [64], Tie et al. [215], PharmacyGPT [143], Zhang et al. [275], Fu et al. [62], Mental-LLM [248], Peters and Matz [173]
		PaLM	Liu et al. [139]
FLAN-T5		Mental-LLM [248]	
Llama/Alpaca/Vicuna		Tie et al. [215], Mental-LLM [248]	
Simulator/Agent	Recommendation	GPT family	RecMind [236], InteRecAgent [80], RecAgent [231], Graph-Toolformer [273]
		LaMDA	RecLLM [61]
	Dialogue System	ChatGPT	UGRO [78]
		Llama	Kong et al. [106]
Intelligent Assistant	GPT-3.5	PersonaLLM [90]	
Classifier/Detector	User Profiling	GPT family	Zhang et al. [271], LoT [76], Mu et al. [153], SentimentGPT [104], Ziems et al. [292]
		LLaMA	Mu et al. [153]
	Healthcare	GPT family	Qi et al. [177], Qin et al. [178], ALEX [97]
	Fraud Detection	GPT family	Shukla et al. [200], Zhu et al. [291]
		T5	Spam-T5 [110]
	Discrimination Detection	GPT-3	Chiu et al. [29]
		Llama	Nguyen et al. [156]
		FLAN-T5	Del Arco et al. [40]
	Misinformation Detection	ChatGPT	Li et al. [125]
		Llama	Pavlyshenko [167]
LLM-Gen Text Detection	ChatGPT	Bhattacharjee and Liu [14]	
Scoring Function	Recommendation	ChatGPT	Chat-REC [63], Liu et al. [133], Kang et al. [100], BookGPT [288], Dai et al. [36]
		LaMDA	RecLLM [61]
		Llama	TallRec [8]
	Dialogue System	ChatGPT	Hu et al. [78]
	Misinformation Detection	ChatGPT	Yang et al. [251]
Explainer	Recommendation	ChatGPT	Chat-REC [63], Liu et al. [133],
	Discrimination Detection	GPT family	Wang et al. [227], Ziems et al. [292]
	Education	Llama/Vicuna	Bao et al. [9]
	Healthcare	GPT-3.5	ALEX [97]
Chatbot	Recommendation	ChatGPT	Chat-REC [63], Lin and Zhang [129], He et al. [72], GeneRec [233],
	Dialogue System	GPT family	DiagGPT [19], Cho et al. [30], RefGPT [252], Hudeček and Dušek [83], Tu et al. [217], Zheng et al. [286]
	Intelligent Assistant	ChaGPT	Lakkaraju et al. [112], Hassan et al. [68]
	Healthcare	ChaGPT	Chen et al. [23], ChatDoctor [127]
	Education	Llama	EduChat [37]

Table 6: Representative approaches of LLMs-as-Enhancers, LLMs-as-Controllers, LLMs-as-Evaluators.

Paradigms	Roles	Applications	LLM Backbones	Models
LLMs-as-Enhancers	Profiler	Recommendation	ChatGPT	GENRE [135], HKPF [260], KAR [246]
			GPT-3	NIR [229]
			Llama	ONCE [136], PALR [26], Richardson et al. [185]
			ChatGLM	LGIR [46]
	Feature Encoder	Recommendation	GPT family	GPT4SM [171], Li et al. [123]
			Llama family	LKPNR [187], LLM4Jobs [121]
			ChatGLM	LKPNR [187], KAR [246]
			RWKV	LKPNR [187]
		User Profiling	GPT family	SentimentGPT [104]
	Knowledge Augmenter	Recommendation	ChatGPT	HKPR [260], KAR [246], Fang et al. [52]
			GPT-3	MINT [154], LLM-Rec [145]
			GPT-2	GPT4Rec [120], Li et al. [122]
			Alpaca	Acharya et al. [3]
		Dialogue System	GPT-3	TacoBot [150]
		Healthcare	GPT family	PULSAR [118], AugESC [284], Schlegel et al. [196], LLM-PTM [268]
		Discrimination Detection	GPT-3	Cohen et al. [35]
	Data Generator	Intelligent Assistant	OPT-175B	VA-Model [7]
		Recommendation	GPT family	LLM4Jobs [121], Graph-Toolformer [273]
		Healthcare	ChatGPT	Tang et al. [211], Chen et al. [23], Wang et al. [228]
		Misinformation Detection	ChatGPT	Su et al. [203], Leite et al. [114], Pan et al. [164]
Discrimination Detection		ChatGPT	Veselovsky et al. [221]	
User Profiling		PaLM	Deng et al. [41]	
Dialogue System		GPT-4	Foosherian et al. [60]	
Fraud Detection		GPT-4	Yang and Menczer [253], Ayoobi et al. [6]	
	LLM-Gen Text Detection	GPT family	Yu et al. [266], Chen and Shu [21], Liu et al. [138]	
LLMs-as-Controllers	Pipeline Controller	Recommendation	LaMDA	RecLLM [61]
			ChatGPT	Chat-REC [63]
			Vicuna	LLM4Jobs [121]
		Dialogue System	GPT-4	Foosherian et al. [60]
LLMs-as-Evaluators	Evaluator	Dialogue System	GPT family	Svikhnushina and Pu [207], Huynh et al. [84], Zheng et al. [285], LLM-Eval [132]
			Claude	Huynh et al. [84], Zheng et al. [285]
			TNLG, BLOOM, OPT, Flan-T5	Huynh et al. [84]
		Recommendation	GPT family	GIRL [287], Wang et al. [235]
		Education	GPT-3	Bhat et al. [13]