

Demonstration of DB-GPT: Next Generation Data Interaction System Empowered by Large Language Models

Siqiao Xue[◇], Danrui Qi[♣], Caigao Jiang[◇], Wenhui Shi[◇], Fangyin Cheng[♥], Keting Chen[◇],
 Hongjun Yang[◇], Zhiping Zhang[♡], Jianshan He[◇], Hongyang Zhang[◇], Ganglin Wei[◇],
 Wang Zhao, Fan Zhou[◇], Hong Yi, Shaodong Liu[♣], Hongjun Yang[◇], Faqiang Chen^{◇,*}
[◇]Ant Group, [♡]Alibaba Group, [♥]JD Group, [♣]Meituan,
[♣]Southwestern University of Finance and Economics, China,
[♣]Simon Fraser University, Canada

ABSTRACT

The recent breakthroughs in large language models (LLMs) are positioned to transition many areas of software. The technologies of interacting with data particularly have an important entanglement with LLMs as efficient and intuitive data interactions are paramount. In this paper, we present DB-GPT, a revolutionary and product-ready Python library that integrates LLMs into traditional data interaction tasks to enhance user experience and accessibility. DB-GPT is designed to understand data interaction tasks described by natural language and provide context-aware responses powered by LLMs, making it an indispensable tool for users ranging from novice to expert. Its system design supports deployment across local, distributed, and cloud environments. Beyond handling basic data interaction tasks like Text-to-SQL with LLMs, it can handle complex tasks like generative data analysis through a Multi-Agents framework and the Agentic Workflow Expression Language (AWEL). The Service-oriented Multi-model Management Framework (SMMF) ensures data privacy and security, enabling users to employ DB-GPT with private LLMs. Additionally, DB-GPT offers a series of product-ready features designed to enable users to integrate DB-GPT within their product environments easily. The code of DB-GPT is available at Github¹ which already has **over 10.7k stars**. Please install DB-GPT for your own usage with the instructions² and watch a 5-minute introduction video on Youtube³ to further investigate DB-GPT.

1 INTRODUCTION

Large language models (LLMs) such as ChatGPT and GPT-4 have showcased their remarkable capabilities in engaging in human-like communication and understanding complex queries, bringing a trend of incorporating LLMs in various fields. Data interaction, which aims to let users engage with and understand their data, enabling the retrieval, analysis, manipulation, and visualization of data to derive insights or make decisions. In the realm of interacting with data, LLMs pave the way for natural language interfaces, enabling users to express their data interaction tasks through natural language and leading to more natural and intuitive data interactions.

Nonetheless, how to enhance the data interaction tasks with LLMs to provide users reliable understanding and insights to their data still remains an open question. One straightforward approach is to

directly provide commonly used LLMs, such as GPT-4, with instructions on how to interact via few-shot prompting or in-context learning. Moreover, to further facilitate the intelligent interactions with data, many works [1, 8, 14] have incorporated the LLM-powered automated reasoning and decision process (a.k.a., multi-agents frameworks) into the data interaction process. However, these multi-agents frameworks are usually task-specific instead of task-agnostic, limiting their usage to a broad range of tasks. Meanwhile, the interaction with data includes a variety of tasks in practice. For example, it includes the Text-to-SQL / SQL-to-Text tasks, the generation of data analytics, the generation of enterprise report analysis and business insights, etc. It is necessary for users to arrange the workflow of multi-agents according to their own needs. The existing effort [1] does not consider abundant data interaction needs. Finally, though being important, the privacy-sensitive setup for LLM-empowered data interaction is under-investigated. The previous efforts [3, 9] are not designed for data interaction tasks.

To overcome these limitations, our key idea is to propose an open-sourced Python library *DB-GPT* supporting data interaction by using multi-agents with flexible arrangement and privacy-sensitive setup. This idea, however, introduces three main challenges, the first challenge (**C1**) is the design of multi-agents framework for supporting database interaction. The second challenge (**C2**) is the declarative expression supporting arrange multi-agents flexibly. The third challenge (**C3**) focuses on the design of private LLM-empowered data interaction.

To solve C1, we propose the Multi-Agents framework in *DB-GPT* which automates the database interaction tasks. Once users have entered their final goals, the Multi-Agents framework can free their hands, autonomously generate the planning of tasks and execute particular tasks. **To solve C2**, we propose a declarative language called *Agentic Workflow Expression Language (AWEL)* in *DB-GPT*. With AWEL, users can implement their execution plan for multi-agents with simple expression (i.e. few lines of code). Furthermore, to make users more code-free, *DB-GPT* also provides an interface for users constructing their Agentic Workflow with only drag and drop. **To solve C3**, we propose *Service-oriented Multi-model Management Framework (SMMF)* in *DB-GPT* to support users to run *DB-GPT* with their private LLMs in their own execution environment. All the interactions among users, LLMs and data are performed locally, which definitely promises users' privacy.

¹<https://github.com/eosporos-ai/DB-GPT>

²<https://github.com/eosporos-ai/DB-GPT#install>

³https://youtu.be/n_8R11ENy14

Table 1: Comparison between DB-GPT and other tools.

		LangChain [1]	LlamaIndex [8]	PrivateGPT [9]	ChatDB [4]	DB-GPT
System Components	Multi-Agents Framework	✓	✓	✗	✗	✓
	Multi-LLMs Support	✓	✓	✗	✓	✓
	RAG from Multiple Data Sources	✓	✓	✗	✗	✓
	Agent Workflow Expression Language	✗	✗	✗	✗	✓
	Fine-tuned Text-to-SQL Model	✗	✓	✗	✗	✓
Data Interaction Functionalities	Text-to-SQL / SQL-to-Text	✓	✓	✗	✓	✓
	Chat2DB / Chat2Data / Chat2Excel	✓	✓	✗	✓	✓
	Data Privacy and Security	✗	✗	✓	✗	✓
	Multilingual Interactions	✗	✗	✗	✓	✓
	Generative Data Analysis	✗	✗	✗	✗	✓




Figure 1: System Design of DB-GPT

Additionally, the *DB-GPT* community extends its support beyond basic functionalities, offering a suite of product-ready features designed to enhance data interaction capabilities. These include advanced knowledge extraction from diverse data sources for more accurate answers to users’ queries, specialized fine-tuning of Text-to-SQL Large Language Models (LLMs) to facilitate seamless database queries, and a user-friendly front-end interface for more convenient interaction. Furthermore, *DB-GPT* supports multilingual functionality, accommodating both English and Chinese, thereby broadening its applicability and ease of use across different linguistic contexts. With these comprehensive, product-ready considerations, *DB-GPT* is equipped to handle intricate data interaction tasks, such as generative data analysis, enabling users to seamlessly integrate and leverage its powerful functionalities within their product environments. This holistic approach ensures that *DB-GPT* is not just a library, but a complete solution for developers and businesses aiming to harness the full potential of AI in the process of interacting with data. Table 1 shows the comparison between *DB-GPT* and other popular tools from two main perspectives: system components and data interaction functionalities, showing the superiority of *DB-GPT*.

To summarize, we make the following contributions: 1) we propose *DB-GPT*, an open-sourced and product-ready library supporting an end-to-end interaction with data. 2) we propose Multi-Agents Framework in *DB-GPT* for solving complex data interaction tasks like generative data analysis. 3) we propose *Agentic Workflow Expression Language (AWEL)* to enhance the practicability and flexibility of Multi-Agents in *DB-GPT*. 4) we propose *Service-oriented Multi-model Management Framework (SMMF)* to promise the users’ privacy from the model perspective in *DB-GPT*. 5) we deploy *DB-GPT* as an application with user-friendly interface and demonstrate its utility. We also open-sourced the implementation of *DB-GPT* on [Github](#), which already has **over 10.7k stars**.

2 SYSTEM DESIGN

The overall system design of *DB-GPT* is depicted in Figure 1. *DB-GPT* includes four layers, i.e. the protocol layer, the module layer, the server layer and the application layer. In this section, we delineate the design of each phase with a top-down manner. There are also other layers making *DB-GPT* product-ready. We also introduce the design of these layers in this section.

2.1 The Application Layer

The application layer encompasses the array of data interaction functionalities supported by *DB-GPT*. These include, but are not limited to, Text-to-SQL/SQL-to-Text, chat-to-database interactions (chat2db), chat-to-data queries (chat2data), chat-to-Excel operations (chat2excel), chat-to-visualization commands (chat2visualization), generative data analysis, and question answering based on knowledge bases. These functionalities include the majority of foundational tasks associated with data interaction, illustrating the comprehensive capabilities of the *DB-GPT* framework.

2.2 The Server Layer

The server layer in *DB-GPT* is an optional component that manages external inputs, such as HTTP requests, by integrating them with domain knowledge to guide lower-tier layers, i.e. the Module Layer. This layer’s optional status allows for direct communication between the application layer and the module layer in simple scenarios. In contexts that necessitate external inputs, the server layer acts as a supplementary intermediary, underscoring its utility in supporting a wider range of applications.

2.3 The Module Layer

The module layer of *DB-GPT* is composed by *Service-oriented Multi-model Management Framework (SMMF)*, *Retrieval-Augmented Generation (RAG) from Multiple Data Source* and *Multi-Agents Framework*. The three parts of the module layer are most important to support users' interaction with their data, which is shown in the Application Layer.

Service-oriented Multi-model Management Framework (SMMF). The Service-oriented Multi-model Framework (SMMF) in the context of *DB-GPT* aims at facilitating model adaptation, enhancing deployment efficiency, and optimizing performance. SMMF offers a streamlined platform for the deployment and inference of Multi-Large Language Models (Multi-LLMs), enabling local execution of users' own LLMs to ensure data privacy and security.

SMMF is underpinned by two core components: the model inference layer and the model deployment layer. The inference layer supports various LLM inference frameworks, enhancing the framework's flexibility. The deployment layer connects inference mechanisms with model serving capabilities, incorporating an API server and a model handler for robust functionality. At its core, the model controller manages metadata, integrating the deployment process, while the model worker establishes connectivity with inference and infrastructure, ensuring efficient model operation. Through SMMF, *DB-GPT* provides an efficient approach to deploying machine learning models in a cloud environment, highlighting the framework's potential in improving adaptability, performance, and data security in MaaS applications.

Retrieval-Augmented Generation (RAG) from Multiple Data Source. While LLMs are usually trained on enormous bodies of open sourced or other parties' proprietary data, RAG [7] is a technique for augmenting LLMs' knowledge with additional and often private data. Shown in Figure 2, our RAG pipeline consists of three stages: knowledge construction, knowledge retrieval and adaptive In-Contextual Learning (ICL) [2] strategies.

For knowledge construction, *DB-GPT* constructs a knowledge base according to multiple data sources provided by users. Contents in each data source are segmented into paragraphs, with each paragraph encoded into a multidimensional vector using a neural encoder. Notably, *DB-GPT* enhances traditional vector-based knowledge representation by integrating inverted index and graph index methods, facilitating precise context-relevant data retrieval. For knowledge retrieval, upon receiving a query x , it is transformed into a query vector q . *DB-GPT* then identifies the top- k paragraphs within the knowledge base that are most relevant to q . *DB-GPT* employs diverse retrieval strategies for prioritizing relevant documents, including ordering based on the cosine similarity of their embedded vectors, as well as categorization according to keyword similarity. In the adaptive iterative contextualization phase, *DB-GPT* employs Interactive Contextual Learning (ICL) for generating responses. ICL enhances *DB-GPT*'s response by integrating knowledge retrieval results during LLMs' inference. It incorporates them into a predefined prompt template to get response from LLM. The efficacy of ICL depends on specific configurations such as prompt templates. Our *DB-GPT* system provides various strategies for prompt formulation and incorporates privacy measures to protect private information. Due to the page limit, please see [17] for the full details.

Multi-Agents Framework. Inspired by [MetaGPT](#) and [AutoGen](#), when dealing with challenging data interaction tasks such as generative data analysis, *DB-GPT* proposes its own Multi-Agent framework. The proposed framework leverages the specialized capabilities and communicative interactions of multiple agents to effectively address multifaceted challenges. For example, consider the task of constructing detailed sales reports from at least three distinct dimensions. The Multi-Agent framework initiates this process by deploying a planning agent to devise a comprehensive strategy, which includes the creation of: 1) a donut chart for the analysis of total sales by product category, 2) a bar chart for examining sales data from the perspective of user demographics, and 3) an area chart for evaluating monthly sales trends. Subsequent to the planning phase, dedicated chart-generating agents are tasked with the production of these visual representations, which are then aggregated by the planner and presented to users.

Compared to [MetaGPT](#) and [AutoGen](#), *DB-GPT*'s Multi-Agent framework archives the entire communication history among its agents within a local storage system, thereby significantly enhancing the reliability of the generated content of agents. Furthermore, in contrast to the [LlamaIndex](#) framework, which prescribes a set of constrained behaviours tailored to specific use cases, *DB-GPT*'s framework offers flexibility which allows users to custom-define agents tailored to their specific data interaction tasks, thus affording a broader applicability across various domains.

2.4 The Protocol Layer

The protocol layer in *DB-GPT* mainly includes *Agentic Workflow Expression Language (AWEL)*, which adopts the big data processing concepts of [Apache Airflow](#). By leveraging Directed Acyclic Graphs (DAGs), AWEL orchestrates workflows, aligning with Apache Airflow's mission to efficiently define, schedule, and oversee complex data pipelines and workflows.

In Apache Airflow, the core components of these workflows are operators, where each operator represents a discrete task or operation capable of executing defined actions. Reflecting this approach, *DB-GPT*'s AWEL models each agent as a distinct operator, thus enabling users to intricately design their agent-based workflows. This is achieved by interconnecting multiple agents to construct a DAG. Such a design grants users the flexibility to manipulate the flow of information between agents. Consequently, users can seamlessly integrate their comprehension of specific data interaction tasks with the actionable insights generated by LLM-based agents.

Employing AWEL within *DB-GPT* empowers it to support a variety of tasks including stream processing, batch processing, and asynchronous operations. This capability significantly bolsters *DB-GPT*'s effectiveness and applicability in navigating the complexities of real-world production environments.

2.5 Other Layers

Visualization Layer. The visualization layer aims to display the answers returned by *DB-GPT* to the users with elegance. For scenarios involving purely textual question-and-answer formats, this layer exhibits the textual outputs generated by *DB-GPT*. When the tasks necessitate the generation of charts, *DB-GPT* renders these charts within its front-end, facilitating user interaction with the displayed

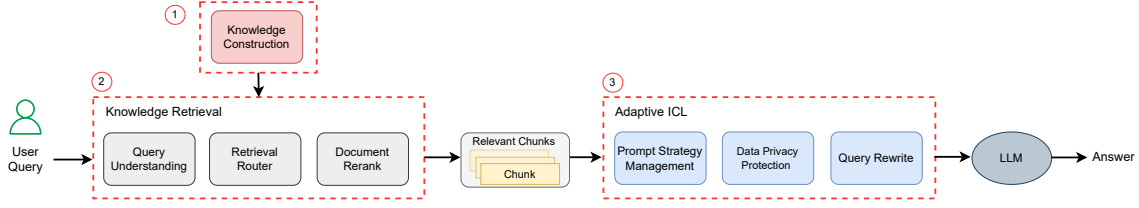


Figure 2: The RAG architecture in DB-GPT

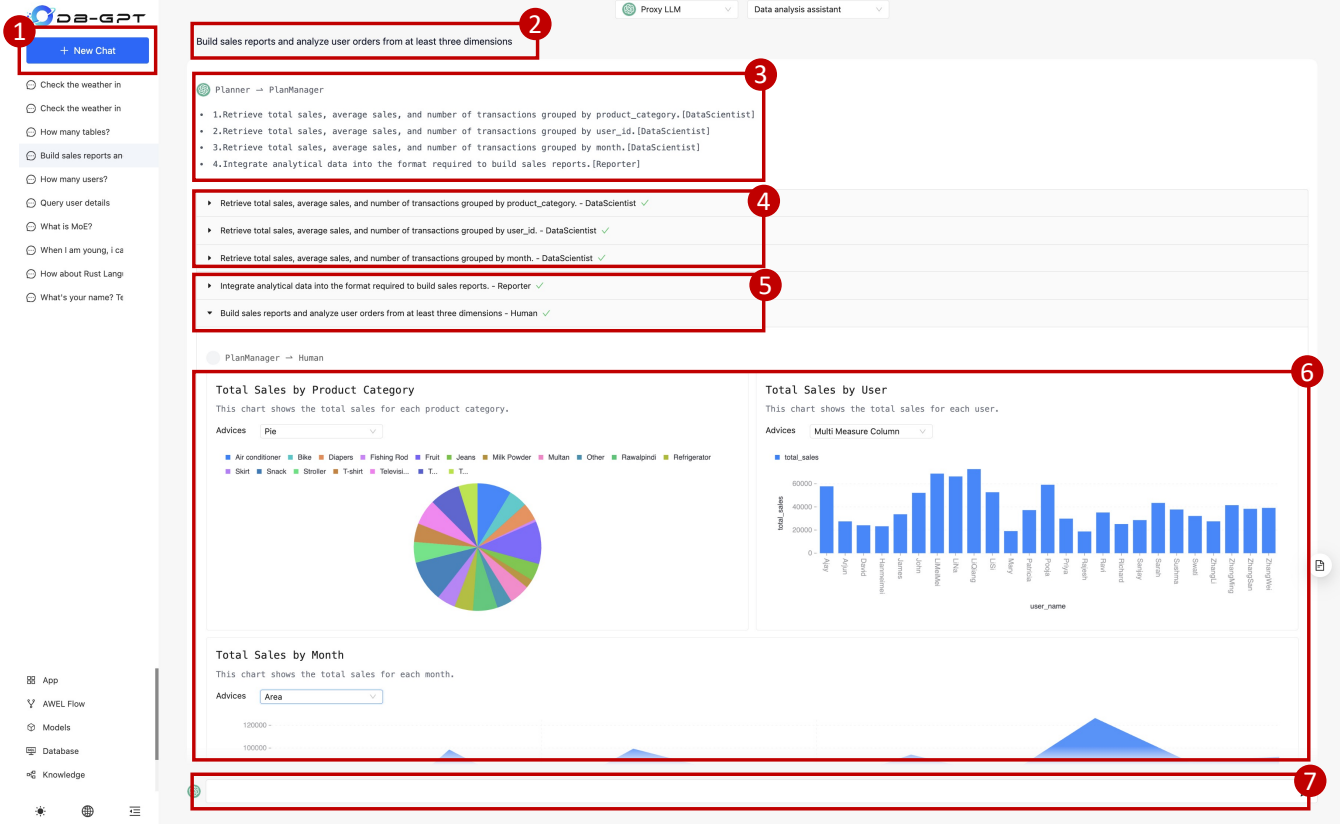


Figure 3: Demonstration of DB-GPT

charts. The significance of possessing a sophisticated visualization layer within DB-GPT cannot be overstated, as practical applications demand that users engage in multiple interactions with their data to successfully complete their data interaction tasks. Superior visualization capabilities significantly expedite users' comprehension of their data, thereby enhancing the overall effectiveness and user experience of the system.

Text-to-SQL Fine-Tuning. Although LLMs, e.g., CodeX and ChatGPT, have shown successful results for Text-to-SQL, they still have a gap with the fine-tuned alternatives in specific application scenarios. Consequently, tailoring LLMs to domain-specific Text-to-SQL datasets emerges as a crucial step towards enhancing their comprehension of prompts and facilitating superior outcomes. Within DB-GPT, we have introduced a component, termed DB-GPT-Hub, which serves to encapsulate the Text-to-SQL fine-tuning process.

This module enables users to refine their Text-to-SQL LLMs using publicly available LLMs hosted on Huggingface in conjunction with their own Text-to-SQL data pairs. Moreover, our SMMF framework accords users the flexibility to employ their fine-tuned LLMs in a localized manner, which underscores our commitment to data privacy and security.

Execution Environments. DB-GPT is capable of operating within distributed environments through the employment of the distributed framework Ray, as well as within cloud ecosystems such as AWS Cloud, Alibaba Cloud, and private cloud configurations maintained by users. This operational flexibility underscores DB-GPT's adeptness at accommodating a variety of data storage contexts. Furthermore, it exemplifies the wide applicability of DB-GPT across diverse operational requirements and environments.

3 DEMONSTRATION

The demonstration setup includes a table need to be standardized and a laptop. The laptop must connect to the Internet for visitors can use DB-GPT smoothly with OpenAI's GPT service. Visitors can also choose local models such as Qwen and GLM. If the conference Internet fails, a mobile hotspot (established via cell phone) can also be used for running DB-GPT.

Figure 3 illustrates the capability of DB-GPT to perform generative data analysis. When users are faced with a data interaction task, they initiate the process by starting a new chat session (area ①) and entering a command such as "Build sales reports and analyze user orders from at least three distinct dimensions" (area ②). DB-GPT undertakes this task utilizing its Multi-Agent framework, which begins with invoking a planner to generate a four-step strategy tailored to the task (area ③). Then, three specialized agents, designated for the creation of data analytics charts, proceed to generate sales reports (area ④). These report takes into account various dimensions, including product category, user name and month. Another agent, dedicated to aggregating these charts, collects, organizes, and presents them on the front-end interface (area ⑤). The interface allows users to interact with the displayed charts, offering the flexibility to alter chart types according to their preferences (area ⑥). If users need further data interaction tasks to be performed, they can continue to engage with their data through natural language inputs (area ⑦).

4 CONCLUSION

In this paper, we proposed DB-GPT, a revolutionary and product-ready Python library that understands data interaction tasks described by natural language and provides responses powered by LLMs. With the four-layer system design, DB-GPT can handle complex data interaction tasks with privacy consideration. In the future, DB-GPT will adapt more data interaction needs with its code-free agentic workflow.

Future research directions include: 1) introducing powerful agents providing more powerful abilities, such as time series predictions [6, 15, 19–21] based on historical data and predictive decision abilities [13, 18] and automatic data preparation [10–12]; 2) the integration of more model training techniques. In addition to pre-training, the community is also interested in continual learning techniques for language models, such as continual pre-training [5], prompt learning [16, 22].

REFERENCES

- [1] Harrison Chase. 2022. LangChain. <https://github.com/hwchase17/langchain>
- [2] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A Survey on In-context Learning. <https://api.semanticscholar.org/CorpusID:255372865>
- [3] H2O.ai. 2023. *H2OGPT*. <https://github.com/h2oai/h2ogpt>
- [4] Chenxu Hu, Jie Fu, Chenzhuang Du, Simian Luo, Junbo Zhao, and Hang Zhao. 2023. ChatDB: Augmenting LLMs with Databases as Their Symbolic Memory. *arXiv:2306.03901* [cs.AI]
- [5] Gangwei Jiang, Caigao Jiang, Siqiao Xue, James Y. Zhang, Jun Zhou, Defu Lian, and Ying Wei. 2023. Towards Anytime Fine-tuning: Continually Pre-trained Language Models with Hypernetwork Prompt. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [6] Ming Jin, Qingsong Wen, Yuxuan Liang, Chaoli Zhang, Siqiao Xue, Xue Wang, James Zhang, Yi Wang, Haifeng Chen, Xiaoli Li, Shirui Pan, Vincent S. Tseng, Yu Zheng, Lei Chen, and Hui Xiong. 2023. Large Models for Time Series and Spatio-Temporal Data: A Survey and Outlook. *arXiv:2310.10196* [cs.LG]
- [7] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kuttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *ArXiv abs/2005.11401* (2020). <https://api.semanticscholar.org/CorpusID:218869575>
- [8] Jerry Liu. 2022. *LlamaIndex*. <https://doi.org/10.5281/zenodo.1234>
- [9] Iván Martínez, Daniel Gallego Vico, and Pablo Orgaz. 2023. *PrivateGPT*. <https://github.com/imartinez/privateGPT>
- [10] Danrui Qi, Jinglin Peng, Yongjun He, and Jiannan Wang. 2023. Auto-FP: An Experimental Study of Automated Feature Preprocessing for Tabular Data. *arXiv:2310.02540* [cs.LG]
- [11] Danrui Qi and Jiannan Wang. 2024. CleanAgent: Automating Data Standardization with LLM-based Agents. *arXiv:2403.08291* [cs.LG]
- [12] Danrui Qi, Weiling Zheng, and Jiannan Wang. 2024. FeatAug: Automatic Feature Augmentation From One-to-Many Relationship Tables. *arXiv:2403.06367* [cs.LG]
- [13] Chao Qu, Xiaoyu Tan, Siqiao Xue, Xiaoming Shi, James Zhang, and Hongyuan Mei. 2023. Bellman Meets Hawkes: Model-Based Reinforcement Learning via Temporal Point Processes. <https://arxiv.org/abs/2201.12569>
- [14] Toran Bruce Richards. 2022. AutoGPT. <https://github.com/Significant-Gravitas/AutoGPT>
- [15] Xiaoming Shi, Siqiao Xue, Kangrui Wang, Fan Zhou, James Y. Zhang, Jun Zhou, Chenhao Tan, and Hongyuan Mei. 2023. Language Models Can Improve Event Prediction by Few-Shot Abductive Reasoning. In *Advances in Neural Information Processing Systems*. <https://arxiv.org/abs/2305.16646>
- [16] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. 2022. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 139–149.
- [17] Siqiao Xue, Caigao Jiang, Wenhui Shi, Fangyin Cheng, Keting Chen, Hongjun Yang, Zhiping Zhang, Jianshan He, Hongyang Zhang, Ganglin Wei, Wang Zhao, Fan Zhou, Danrui Qi, Hong Yi, Shaodong Liu, and Faqiang Chen. 2023. DB-GPT: Empowering Database Interactions with Private Large Language Models. *arXiv preprint arXiv:2312.17449* (2023). <https://arxiv.org/abs/2312.17449>
- [18] Siqiao Xue, Chao Qu, Xiaoming Shi, Cong Liao, Shiyi Zhu, Xiaoyu Tan, Lintao Ma, Shiyu Wang, Shijun Wang, Yun Hu, Lei Lei, Yangfei Zheng, Jianguo Li, and James Zhang. 2022. A Meta Reinforcement Learning Approach for Predictive Autoscaling in the Cloud. In *KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 14 - 18, 2022*, Aidong Zhang and Huzefa Rangwala (Eds.). ACM, 4290–4299. <https://doi.org/10.1145/3534678.3539063>
- [19] Siqiao Xue, Xiaoming Shi, Zhixuan Chu, Yan Wang, Fan Zhou, Hongyan Hao, Caigao Jiang, Chen Pan, Yi Xu, James Y. Zhang, Qingsong Wen, Jun Zhou, and Hongyuan Mei. 2023. EasyTPP: Towards Open Benchmarking the Temporal Point Processes. *arXiv preprint arXiv:2307.08097* (2023). <https://arxiv.org/abs/2307.08097>
- [20] Siqiao Xue, Xiaoming Shi, Hongyan Hao, Lintao Ma, James Zhang, Shiyu Wang, and Shijun Wang. 2021. A Graph Regularized Point Process Model For Event Propagation Sequence. In *IJCNN*. 1–7.
- [21] Siqiao Xue, Xiaoming Shi, Y James Zhang, and Hongyuan Mei. 2022. HYPRO: A Hybridly Normalized Probabilistic Model for Long-Horizon Prediction of Event Sequences. In *Advances in Neural Information Processing Systems*. <https://arxiv.org/abs/2210.01753>
- [22] Siqiao Xue, Yan Wang, Zhixuan Chu, Xiaoming Shi, Caigao Jiang, Hongyan Hao, Gangwei Jiang, Xiaoyun Feng, James Zhang, and Jun Zhou. 2023. Prompt-augmented Temporal Point Process for Streaming Event Sequence. In *NeurIPS*.