

Unstructured and structured data: Can we have the best of both worlds with large language models?

Wang-Chiew Tan
Reality Labs Research, Meta

1 Introduction

We are witnessing rapid advancements in the area of large language models (LLMs). A search on Google Scholar shows there are about 3,910 papers with “large language models” in the titles of the papers in 2022. As of April 12, 2023, there are already 1700 articles with LLMs in their titles. In addition to Google Scholar, we are also witnessing huge volumes of blog posts, news articles, twitter feeds, and open-source repositories around LLMs that have sprung up in recent months.

Perhaps ChatGPT (released on November 30, 2022) is the epitome of this LLM revolution that truly unleashed and showcased, to the masses, the power of what has been brewing in the natural language and machine learning communities in recent years. There are many things ChatGPT¹ can do and does so impressively by generating intelligent human-like responses to your questions. Through natural language as input and output², it can solve non-trivial mathematical problems, to a certain extent [? ?], translate your specification into code in a programming language of your choice (e.g., [? ?]), help you write proeses in different styles and the list of accolades goes on [? ?].

In addition to its ability to answer questions, one can also prompt it with tables, such as CSV files, and ask questions over the tables in natural language. More interestingly, it is also possible to prompt ChatGPT with both text and tables and ask questions over the two types of information seamlessly. A little more perseverance in this exercise quickly reveals that ChatGPT has a limit on how much one can input with each prompt and how much information it will retain, at least based on my experience when I tried this at the end of March 2023.

It is natural to wonder whether we can use ChatGPT, with some extensions, as a system for storing and querying data, with natural language as the primary medium of input and output, which it excels at. What are the challenges of doing so, and how can we, as database practitioners and theoreticians, make progress in this context?

2 Unstructured and structured data

A lot of data, including text, images, audio, and videos, sits “outside the box” today. Often, such unstructured data contain multiple modalities simultaneously. For example, we often find text in images [? ?], and we may also find text associated with videos and/or images on the web. Unstructured data is prevalent to a large extent because it is easily authored and shared by users [? ?] through a variety of apps and authoring tools that are widely available. Such data is often queried with keyword search and today, they can also be queried in natural language with LLMs. Typically the cost of devising a schema and setting up a database for querying the data inhibits the use of a database management system (DBMS) upfront. However, as data scales, the need for structure and semantics becomes more critical, so as to enable faster and more accurate retrieval of content. For example, organizing photos by year, trips, or entity types (such as people or pets) adds some structure, which makes answering certain types of queries much faster and more accurate. However, answering complex queries such as “*when was the last time I went to the coffee shop beside restaurant Italio?*” or “*how many times did I celebrate Anna’s birthday with*

¹I use the term “ChatGPT” as a representative for general chat systems based on LLMs for the rest of this article.

²GPT-4 can accept image and text inputs.

a mango cake?” requires non-trivial reasoning and computation over the data that goes beyond the capabilities of LLMs today [? ?].

Enterprise data sits on the other end of the spectrum. It is, for the most part, not authored by everyday users and is highly structured, often sitting "in a box" in some DBMS. Enterprise data comes with a well-devised schema and is typically highly optimized to serve a sizable query workload with great efficiency. Such data is often queried directly with SQL which is adequately expressive for specifying complex queries such as those with aggregates and recursion. Significant research has also been carried out to enable querying a database with natural language (e.g., [? ?]), where questions are posed in natural language and translated into SQL, which can then be executed over the DBMS. A DBMS is highly optimized to handle large amounts of data and can also perform transactions with ACID guarantees [?]. However, the core dbms does not understand natural language and it tends to fall short in its ability to query unstructured data, which are often stored as blobs and adding semantics to the blobs require additional effort. For example, “*show me the sales numbers over Black Friday and Cyber Monday last year*” requires commonsense knowledge on what “sales numbers” mean and how that maps to relevant table attribute(s), which may be stored under different names in different databases, and when Black Friday and Cyber Monday occurred. Another example is “*find all items with good reviews that are similar to these images,*” which requires matching images (semantically) and interpreting what “good reviews” mean based on the data.

Despite the divide in how unstructured and structured data are managed today, the desire to query in natural language is common to both. At the same time, it is unreasonable and unnatural to expect all unstructured data to fit in some structure, or vice versa to leverage one system for querying. So, how can we effectively query both types of data with the help of LLMs, which possess tremendous knowledge and language understanding?

3 The best of both worlds with LLMs?

LLMs contain tremendous parametric knowledge in their model parameters but lack the ability to incorporate external data (i.e., data outside their model). Hence, if a model is trained based on data up to, say, 2021, it will not provide correct answers about events or facts that require knowledge after 2021. For example, if someone asks the question “*How were the midterm election results of 2022?*” on the webpage with `text-davinci-003` [?], the answer returned is “*The midterm election results of 2022 are not yet available, as the election has not yet taken place.*” This is because `text-davinci-003` is trained with data up to June 2021.

Retrieval-augmented language models (e.g., [? ? ?]) overcome this limitation by adopting a semi-parametric approach to answering queries. They use external data, by first retrieving relevant data from an external data store, and then attempts to answer a question conditioned on the retrieved data and with their parametric knowledge. However, as demonstrated by (Table 5, [?]), such systems can still perform poorly on complex queries involving aggregates and certain types of temporal queries. This is because, for such queries, oftentimes, it is impossible to fit all necessary data for answering the question into the finite-length token input imposed by language models. However, as LLMs get even larger or as more advances are made to increase the token limits imposed by language models, one can anticipate that LLMs will take larger and larger inputs in future and the finite-length token limit may no longer be an issue soon. At the same time, it is likely that there will be even larger datasets to manage and an even larger set of data is required for computing the right answers. So this problem will persist, at least for a while.

A proposal to overcome the above limitation for some types of queries is described in [?]. The paper presents a vision of using views (e.g., as tables) to structure portions of the underlying data sources. The data sources may be of different modalities, such as text, images, videos, or even tables. Views are used to surface important properties about data or associations between data of different modalities and LLMs are used to translate natural language queries into queries (e.g., SQL in this case) that can be executed over the views whenever possible. With this proposal, a key question is to understand when a natural language query can be answered with views. If

a query cannot be answered using views, the system falls back to retrieval-augmented language models to answer the query to its best effort. Alternatively, the two components (view-based and retrieval-based query answering components) can also collaborate to produce a final answer. There are several other questions raised in the paper, such as what views should be materialized? How can one automatically select the “right” views to materialize given an anticipated query workload? And how can one decide when a question is better answered with views, the retrieval system, or even both?

In addition to the above, I will highlight below what I believe are some of the more pertinent questions that may be of immediate interest to the database community:

Query answering with different resources and budget constraints The topic of finding a good query plan to answer a given query was already discussed in [?]. The core of that system relies on two components — views and a retrieval-augmented language model for answering queries. It is conceivable to augment the system with additional components, such as one that generates images given a natural language prompt, which may sometimes be useful for answering certain queries³. A key question then becomes how do we understand when to leverage which component for answering a query or have the components collaborate to derive an answer? Furthermore, LLMs are compute-intensive and can be slow in generating an answer. Some LLMs are also not free. In addition to deciding how best to answer a query with all the available resources, how can one account for the strengths and limitations of each component to enumerate and compare plans for computing an answer to a given query and/or under a given budget? Can we also use a language model to generate a query plan or some parts of it, similar to how language models have been used to self-reason a sequence of steps to derive answers from questions (e.g., [?])?

Provenance Provenance is well-studied for certain classes of SQL queries, and is roughly defined as the source tuples that explains why a tuple is in the result of the query. As mentioned in [?], one should attempt to answer queries using table views whenever possible by translating the natural language query into a SQL query that can be executed over the views. This way, it is possible that provenance can be obtained “for free”. However, SQL queries that are generated by LLMs can be complex, for example, with nested SQL queries and/or aggregates in the FROM or WHERE clauses. For such cases, can we decompose the generated SQL query into a sequence of one or more “simpler queries” instead, where the provenance for simpler queries is well-understood and can be derived easily? If this is not possible, can we strategize a plan for answering the query in a different way so that provenance can be derived? The problem of finding an alternative query plan is related to the discussion in the earlier paragraph, but here the focus is on deriving a plan with sufficiently simple steps to enable provenance.

Retrieval-augmented language models also provide more guarantees for providing evidence for their answers. We are also beginning to witness implementations where sources of answers presented by retrieval-augmented language models are returned as part of the answers [?]. However, more research needs to be carried out to attribute provenance to training data (e.g., [? ?]) to form a more comprehensive picture of provenance for the provided answer.

Prompt Engineering Analysis LLMs have limits on the number of tokens they can take as input. Even if one takes advantage of all the tokens one can use, prompting them with more information does not always translate to better answers as sometimes, presenting the language model with more information confuses the language model. This means one needs to be judicious in what we send to a language model for it to derive a correct answer of high quality⁴.

The answers returned by language models are also sensitive to how they are prompted. Sometimes the same question phrased slightly differently will result in completely different answers. In prompt engineering, the goal is to find the best prompt for the task at hand.

Given the maximum token limit of language models, can we optimize the answers returned by a language model by strategically summarizing relevant data and/or removing irrelevant data? For example, the entity

³As they say, a picture is worth a thousand words.

⁴There are varied ways to answer a question correctly. Some are better than others.

matching system of [?], which uses a language model, immediately performs entity matching more accurately when text descriptions of data are strategically shortened using a simple trick; by keeping only tokens of value, with high TD/IDF. On a more theoretical side, given a suitable definition of what is a prompt and assumptions about language models, can we characterize what types of queries that require access to external data can be answered with one prompt, a finite number of prompts, or an asymptotic number of prompts under a budget of tokens?

4 Conclusion

The field of LLMs is moving fast, both in research and industry. In addition to [?], [?] have also described the challenges of answering queries, in the context of augmented language models and data integration, with a single source or by chaining multiple sources. In [?], the authors described how a DBMS can be extended to leverage LLMs to improve query answering and also pointed to the direction of a hybrid query answering system involving both a DBMS and LLMs. I believe this is only the beginning and we will see many more visions, research, and implementations soon in this area of query answering systems that embrace both unstructured and structured data through the use of large language models.

Acknowledgements Many of the ideas above are inspired from discussions with Alon Halevy and Yuliang Li. I also thank many of my colleagues at Meta — Lambert Mathias, Richard Newcombe, and Luna Dong — for active discussions around large language models from which I have learnt lots and also to Lucian Popa for his feedback on this article.

References

- [1] Chirag Agarwal, Owen Queen, Himabindu Lakkaraju, and Marinka Zitnik. Evaluating explainability for graph neural networks. *Scientific Data*, 10(1):144, 2023.
- [2] Ulrich Aïvodji, Alexandre Bolot, and Sébastien Gams. Model extraction from counterfactual explanations. *arXiv preprint arXiv:2009.01884*, 2020.
- [3] Kenza Amara, Rex Ying, Zitao Zhang, Zhihao Han, Yinan Shan, Ulrik Brandes, Sebastian Schemm, and Ce Zhang. Graphframex: Towards systematic evaluation of explainability methods for graph neural networks. *arXiv preprint arXiv:2206.09677*, 2022.
- [4] Maximilian Augustin, Valentyn Boreiko, Francesco Croce, and Matthias Hein. Diffusion visual counterfactual explanations. *arXiv preprint arXiv:2210.11841*, 2022.
- [5] Steve Azzolin, Antonio Longa, Pietro Barbiero, Pietro Liò, and Andrea Passerini. Global explainability of gnns via logic combination of learned concepts. *arXiv preprint arXiv:2210.07147*, 2022.
- [6] Federico Baldassarre and Hossein Azizpour. Explainability techniques for graph convolutional networks. *CoRR*, abs/1905.13686, 2019.
- [7] Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and R Devon Hjelm. Mine: mutual information neural estimation. *arXiv preprint arXiv:1801.04062*, 2018.
- [8] Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. In *NeurIPS*, 2021.
- [9] Pietro Bongini, Monica Bianchini, and Franco Scarselli. Molecular generative graph neural networks for drug discovery. *Neurocomputing*, 450:242–252, 2021.

- [10] Anshika Chaudhary, Himangi Mittal, and Anuja Arora. Anomaly detection using graph neural networks. In 2019 international conference on machine learning, big data, cloud and parallel computing (COMITCon), pages 346–350. IEEE, 2019.
- [11] Dawei Cheng, Fangzhou Yang, Sheng Xiang, and Jin Liu. Financial time series forecasting with multi-modality graph neural network. Pattern Recognition, 121:108218, 2022.
- [12] Enyan Dai, Tianxiang Zhao, Huaisheng Zhu, Junjie Xu, Zhimeng Guo, Hui Liu, Jiliang Tang, and Suhang Wang. A comprehensive survey on trustworthy graph neural networks: Privacy, robustness, fairness, and explainability. arXiv preprint arXiv:2204.08570, 2022.
- [13] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. Graph neural networks for social recommendation. In The world wide web conference, pages 417–426, 2019.
- [14] Madelyn Glymour, Judea Pearl, and Nicholas P Jewell. Causal inference in statistics: A primer. John Wiley & Sons, 2016.
- [15] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. Strategies for pre-training graph neural networks, 2020.
- [16] Qiang Huang, Makoto Yamada, Yuan Tian, Dinesh Singh, Dawei Yin, and Yi Chang. Graphlime: Local interpretable model explanations for graph neural networks. CoRR, abs/2001.06216, 2020.
- [17] Guillaume Jeanneret, Loïc Simon, and Frédéric Jurie. Diffusion models for counterfactual explanations. In Proceedings of the Asian Conference on Computer Vision, pages 858–876, 2022.
- [18] José Jiménez-Luna, Francesca Grisoni, and Gisbert Schneider. Drug discovery with explainable artificial intelligence. Nature Machine Intelligence, 2(10):573–584, 2020.
- [19] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. CoRR, abs/1609.02907, 2016.
- [20] Thomas N Kipf and Max Welling. Variational graph auto-encoders. arXiv preprint arXiv:1611.07308, 2016.
- [21] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11):2278–2324, 1998.
- [22] Wenqian Li, Yinchuan Li, Zhigang Li, Jianye Hao, and Yan Pang. Dag matters! gflownets enhanced explainer for graph neural networks. arXiv preprint arXiv:2303.02448, 2023.
- [23] Yiqiao Li, Jianlong Zhou, Sunny Verma, and Fang Chen. A survey of explainable graph neural networks: Taxonomy and evaluation metrics. arXiv preprint arXiv:2207.12599, 2022.
- [24] Yiqiao Li, Jianlong Zhou, Boyuan Zheng, and Fang Chen. Ganexplainer: Gan-based graph neural networks explainer. arXiv preprint arXiv:2301.00012, 2022.
- [25] Wanyu Lin, Hao Lan, and Baochun Li. Generative causal explanations for graph neural networks. In International Conference on Machine Learning, pages 6666–6679. PMLR, 2021.
- [26] Wanyu Lin, Hao Lan, Hao Wang, and Baochun Li. Orphicx: A causality-inspired latent variable model for interpreting graph neural networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 13729–13738, 2022.
- [27] Dongsheng Luo, Wei Cheng, Dongkuan Xu, Wenchao Yu, Bo Zong, Haifeng Chen, and Xiang Zhang. Parameterized explainer for graph neural network. In NeurIPS, 2020.
- [28] Jing Ma, Ruocheng Guo, Saumitra Mishra, Aidong Zhang, and Jundong Li. Clear: Generative counterfactual explanations on graphs. arXiv preprint arXiv:2210.08443, 2022.

- [29] Siqi Miao, Mia Liu, and Pan Li. Interpretable and generalizable graph learning via stochastic attention mechanism. In *International Conference on Machine Learning*, pages 15524–15543. PMLR, 2022.
- [30] Tri Minh Nguyen, Thomas P Quinn, Thin Nguyen, and Truyen Tran. Counterfactual explanation with multi-agent reinforcement learning for drug target prediction. *arXiv preprint arXiv:2103.12983*, 2021.
- [31] Phillip E Pope, Soheil Kolouri, Mohammad Rostami, Charles E Martin, and Heiko Hoffmann. Explainability methods for graph convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10772–10781, 2019.
- [32] Phillip E. Pope, Soheil Kolouri, Mohammad Rostami, Charles E. Martin, and Heiko Hoffmann. Explainability methods for graph convolutional neural networks. In *CVPR*, pages 10772–10781, 2019.
- [33] Mario Alfonso Prado-Romero, Bardh Prenkaj, Giovanni Stilo, and Fosca Giannotti. A survey on graph counterfactual explanations: Definitions, methods, evaluation. *arXiv preprint arXiv:2210.12089*, 2022.
- [34] Michael Sejr Schlichtkrull, Nicola De Cao, and Ivan Titov. Interpreting graph neural networks for NLP with differentiable edge masking. *CoRR*, abs/2010.00577, 2020.
- [35] Yunsheng Shi, Zhengjie Huang, Wenjin Wang, Hui Zhong, Shikun Feng, and Yu Sun. Masked label prediction: Unified message passing model for semi-supervised classification. *CoRR*, abs/2009.03509, 2020.
- [36] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *ICML*, volume 70, pages 3319–3328, 2017.
- [37] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks, 2018.
- [38] Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress: Discrete denoising diffusion for graph generation. *arXiv preprint arXiv:2209.14734*, 2022.
- [39] Minh N. Vu and My T. Thai. Pgm-explainer: Probabilistic graphical model explanations for graph neural networks. In *NeurIPS*, 2020.
- [40] Jianian Wang, Sheng Zhang, Yanghua Xiao, and Rui Song. A review on graph neural network methods in financial applications. *arXiv preprint arXiv:2111.15367*, 2021.
- [41] Xiang Wang, Ying-Xin Wu, An Zhang, Xiangnan He, and Tat-Seng Chua. Towards multi-grained explainability for graph neural networks. In *Proceedings of the 35th Conference on Neural Information Processing Systems*, 2021.
- [42] Xiang Wang, Yingxin Wu, An Zhang, Fuli Feng, Xiangnan He, and Tat-Seng Chua. Reinforced causal explainer for graph neural networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2022.
- [43] Xiaoqi Wang and Han-Wei Shen. Gnninterpreter: A probabilistic generative model-level explanation for graph neural networks. *arXiv preprint arXiv:2209.07924*, 2022.
- [44] Dana Warmusley, Alex Waagen, Jiejun Xu, Zhining Liu, and Hanghang Tong. A survey of explainable graph neural networks for cyber malware analysis. In *2022 IEEE International Conference on Big Data (Big Data)*, pages 2932–2939. IEEE, 2022.
- [45] Geemi P Wellawatte, Aditi Seshadri, and Andrew D White. Model agnostic generation of counterfactual explanations for molecules. *Chemical science*, 13(13):3697–3705, 2022.
- [46] Oliver Wieder, Stefan Kohlbacher, Méline Kuenemann, Arthur Garon, Pierre Ducrot, Thomas Seidel, and Thierry Langer. A compact review of molecular property prediction with graph neural networks. *Drug Discovery Today: Technologies*, 37:1–12, 2020.
- [47] Bingzhe Wu, Jintang Li, Junchi Yu, Yatao Bian, Hengtong Zhang, Chaochao Chen, Chengbin Hou, Guoji Fu, Liang Chen, Tingyang Xu, et al. A survey of trustworthy graph learning: Reliability, explainability, and privacy protection. *arXiv preprint arXiv:2205.10014*, 2022.

- [48] Zhitao Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. Gnnexplainer: Generating explanations for graph neural networks. In NeurIPS, pages 9240–9251, 2019.
- [49] Junchi Yu, Tingyang Xu, Yu Rong, Yatao Bian, Junzhou Huang, and Ran He. Graph information bottleneck for subgraph recognition. In International Conference on Learning Representations, 2021.
- [50] Hao Yuan, Jiliang Tang, Xia Hu, and Shuiwang Ji. XGNN: towards model-level explanations of graph neural networks. In Rajesh Gupta, Yan Liu, Jiliang Tang, and B. Aditya Prakash, editors, KDD, pages 430–438, 2020.
- [51] Hao Yuan, Haiyang Yu, Shurui Gui, and Shuiwang Ji. Explainability in graph neural networks: A taxonomic survey. CoRR, 2020.
- [52] Hao Yuan, Haiyang Yu, Jie Wang, Kang Li, and Shuiwang Ji. On explainability of graph neural networks via subgraph explorations. ArXiv, 2021.
- [53] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part I 13, pages 818–833. Springer, 2014.
- [54] He Zhang, Bang Wu, Xingliang Yuan, Shirui Pan, Hanghang Tong, and Jian Pei. Trustworthy graph neural networks: Aspects, methods and trends. arXiv preprint arXiv:2205.07424, 2022.
- [55] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. AI open, 1:57–81, 2020.