© IAEME Publication

THE FUTURE OF BUSINESS INTELLIGENCE: INTEGRATING AI ASSISTANTS LIKE DAX COPILOT INTO ANALYTICAL WORKFLOWS

Tharun Kumar Nallamothu
Senior Software Developer, Microsoft, USA.

ABSTRACT

Artificial intelligence (AI), particularly Nuance's DAX Copilot, is revolutionizing business intelligence (BI) solutions. This AI assistant passively records information by transcribing doctor-patient discussions into structured summaries. A study applied Nuance's domain-aware AI technology to a corporate BI environment, enhancing data quality, reducing effort, and improving decision-making processes. A mixed-methods approach was used, with 20 business analysts using a DAX-enabled BI simulation for 30 days, while a control group used traditional tools. Performance parameters such as report creation time, documentation completeness, and user satisfaction were assessed. "In our business, we pride ourselves on handling reports as efficiently as possible, regardless of market conditions — and the addition of an AI assistant like DAX Copilot directly accelerates our workflow while improving both quality and user experience". Results showed significant improvements, with report creation time reduced by 34%, documentation completeness increased by 28%, and analyst satisfaction increased by 41%. The study recommends adaptable integration strategies and strong privacy controls for promoting AI adoption in BI ecosystems.

1. Introduction

The world of Business Intelligence (BI) has changed over the past ten years. From initially merely descriptive analytics, and dashboard reporting, BI systems are now moving into the direction of predictive and prescriptive analytics, powered by artificial intelligence. One of the most interesting developments on this frontier is the inclusion of AI assistants in analytic flows. These intelligent agents can be used to replace humans in doing repetitive tasks, giving the users a natural language interface to data, and more importantly, to do ambient intelligence, as AI that learns and analyses ongoing user chatter [1].

One of the most well-known examples is Nuance's DAX Copilot (Dragon Ambient Experience Copilot). Originally created to help clinicians by capturing patient encounters automatically, DAX Copilot leverages ambient, conversational and generative AI to serve as an invisible clinical scribe. After Microsoft acquired Nuance, the technology has been expanded into wider clinical and enterprise areas, with products like DAX Express and Dragon Medical One making up the Dragon Copilot suite. No longer are these AI instruments confined to clinical documentation, they're now affecting business workflow and plugging into solutions such as Microsoft Dynamics, Azure AI, and Power BI [2].

On a traditional BI workflow, experts spend a lot of time collecting and formatting data, as well as annotating that data; important work, but also labor-intensive. DAX Copilot-style AI systems should help with much of this mental acrobatics. For example, an AI assistant can record the spoken insights of a team during a strategy meeting, then automatically summarize decisions and recommend which dashboards to use to monitor relevant KPIs. In the same vein, it observes passively the interactions of analysts and creates the records of decisions, queries, and results in a structured format [3].

The potential of combining such ambient AI with BI is massive. It doesn't just ease the pain of documentation, but it also provides better data quality, traceability, and context for development. Analysts can concentrate on analyzing insights rather than writing about them,

resulting in faster and more accurate decisions. In addition, this integration lets there be cognitive connection between more solid data and fluid context, which is typically lost within traditional BI tools [4] [5].

Nuance's DAX Copilot is an AI-powered technology that automates patient encounters, reducing administrative workloads in healthcare. Its capabilities are now being applied to retail, banking, and manufacturing industries, enabling AI assistants to collect and interpret contextual data, perform voice-activated business intelligence queries, automate reporting, and provide on-the-fly analyses. However, businesses must address concerns about data security, organizational readiness, system integration, and user adoption when implementing AI assistants. Proactively addressing these issues is crucial for successful AI adoption and long-term impact [6].

This paper explores how AI assistants like DAX Copilot can be used in business intelligence (BI) systems. It looks at their benefits, potential issues, and how they affect efficiency, user satisfaction, and data quality. A model for using these assistants is proposed. A pilot study compares BI tasks done with and without AI support to measure the value of ambient AI. The findings provide insights for using such AI tools in real-world business settings.

2. Background and Related Work

2.1 Evolution of business intelligence

Business Intelligence (BI) has evolved significantly over the years, particularly with the arrival of computing power, high-capacity storage, cloud tech and user interface. In the old days BI was predominantly a IT/ data engineer/DBA centered area. Older BI technologies were developed on inflexible, highly dependent architectures, which meant that users had to be highly technical to access, manipulate, and analyze data [7] [8].

In the classic BI world, designing queries was all about knowing what the structured query language (SQL) was, how to script, and having an understanding of the schema of the database. Most of the reports are generated centrally and were slow, hampering the ability to make changes and be competitive. Cognos and Analytical outputs Dashboards were more non-interactive (tend to be static) however, such dashboards were primarily targeted towards top management rather than being functional for business users across business functions.

This ecosystem has changed quite a bit with the modern BI platforms like Microsoft Power BI, Tableau, Qlik, etc. These solutions offered 'user-friendly' interfaces, drag-and-drop creation of visualizations for analysis with a self-service approach, allowing a wider range of

users — in particular business people and domain experts — to engage with data without needing to write hard-to-understand code. Cloud-native BI platforms accelerated the ability to share, access, and easily scale in real-time [9].

But despite all these great features there's still a significant challenge: to fully master the advanced features available in these tools for analyzing and modeling. For Power BI, Data Analysis Expressions (DAX) is the foundational formula language that is used to create calculated measures and columns and custom aggregations. Although DAX is awesome and powerful, it is not easy to be trained. DAX is not just about the syntax, it is also about understanding Power BI's tabular data model, thinking in terms of evaluation context (row context and filter context), and using advanced functions [10] [11].

Writing to DAX can be difficult for a lot of a users, many don't have a programming/data science background. This constraint often means that organizations need to leverage only a small cohort of technically capable staffers which themselves form an analytical bottleneck. This gap between business intent and technical execution is still a significant pain point to achieving genuinely democratized BI [12].

Given that organizations are increasingly aspiring to build data-driven cultures, and that they want to empower employees at all layers of the organization to make the best data-based decisions, the need for instruments that diminish this overhead becomes more urgent. This need introduces us to AI-driven assistants—elements that can act as a smart interface between human users and the internal complicated logic of analytical work [13].

2.2 Introducing DAX Copilot

Nuance's DAX Copilot, an AI-powered clinical recording tool, was introduced in early 2023 to automatically capture multi-party interactions during patient visits and convert them into structured clinical summaries. It interacts with EHRs like Epic and MEDITECH Expanse. The fully branded DAX Copilot became available in January 2024, embedding in the Epic EHR system. It allows automated clinical documentation during office and telehealth visits, including encounter recording suggestions and note summarization, reducing administrative duties and allowing clinicians to focus on patient care.

Nuance's DAX Copilot is a revolutionary solution designed to reduce paperwork in healthcare settings. It is part of the Dragon Ambient Experience (DAX) portfolio and uses ambient listening, conversational AI, and large language models (LLMs) to create structured clinical notes, after-visit summaries, and referral letters. It integrates with electronic health record (EHR) systems like Epic and Cerner. DAX Copilot allows clinicians to interact naturally

with the technology, eliminating the need for manual note-taking. It also ensures medical accuracy and compliance with best standards through context-aware document production. The solution is cross-platform compatible with Dragon Medical One and DAX Express solutions, part of the Dragon Copilot suite.

Nuance's DAX Copilot is a clinical documentation interface that uses natural language and ambient AI to allow physicians to speak spontaneously during patient visits. This eliminates the need for physicians to manually record lengthy notes, memorize medical code complexities, or manage complex EHR systems. The AI assistant listens in the background, interprets the context, extracts key clinical details, and automatically generates structured, compliant clinical notes directly into the electronic health record. This allows physicians to have a normal conversation with a patient about symptoms, history, and treatment plans, while the AI assistant interprets the context and extracts key clinical details.

Nuance's DAX Copilot integrates with clinical workflow and EHR systems, ensuring accurate documentation and therapeutic relevance. The assistant is aware of the patient's background, medical history, visit type, and clinical codes. It adjusts documentation output to the clinician's individual encounter and patient profile, ensuring compliance and usefulness in the overall treatment plan [14].

It offers real-time feedback to doctors during the clinical documentation process, suggesting improvements, highlighting missing information, and clarifying as they interact with the document. This not only speeds up note production but also helps professionals learn new EHR workflows faster, as it can identify potential clinical coding difficulties and provide additional details.

It offers not only simplicity but also significant efficiency gains. It reduces the cognitive burden of recalling diagnostic codes, formatting clinical notes, and navigating EHR templates, allowing doctors to operate more fluidly and focus on patient engagement, rather than administrative tasks. This allows healthcare teams to quickly iterate on treatment documentation.

Microsoft and Nuance are integrating AI into their productivity stack, with DAX Copilot being part of the Dragon Ambient Experience ecosystem. This AI-powered tool complements human expertise, allowing medical staff to focus on clinical judgment, advanced patient care, and quality improvement. As AI-powered technologies like DAX Copilot gain traction in healthcare, clinicians will shift from manual documenting to record curation, validation, and

optimization, ensuring thorough, accurate, and clinically useful AI-captured patient stories. The Dragon Ambient Experience ecosystem includes DAX Express and Dragon Medical One.

3. Methodology

The study evaluates the use of AI assistants, specifically a simulated version of Nuance's DAX Copilot, in a business intelligence (BI) work environment. The research uses a mixed-methods approach, involving system implementation and simulation, quantitative evaluation, and qualitative insights. The study focuses on the impact of DAX Copilot on job efficiency, documentation quality, and user acceptance. The methodology includes controlled experiments to measure performance gains and process efficiency, and in-depth interviews to gather user perspectives and feedback.

3.1 Research Design

The study explores the role of AI assistants in enterprise business intelligence (BI) by examining how a simulated version of Nuance's DAX Copilot can improve analytical workflows, documentation quality, and user experience. The research uses an iterative, phased approach, creating a simulated BI environment to measure performance and conducting in-depth interviews to gather context-specific insights. This approach balances measurable outcomes with qualitative knowledge, providing a comprehensive understanding of how an AI assistant designed for clinical recording can be transformed to meet BI professionals' needs.

3.2 Experimental Setup

We modelled a realistic BI environment using Microsoft Power BI, Microsoft Teams, and Microsoft Dynamics 365, and connected it to a prototype assistant that was implemented to simulate the capabilities of DAX Copilot, including background listening, note taking, and summarization. The participants N=30, had at least 2 years of experience in the field of business analytics. Users were split into two groups: a control group used ordinary BI tools, whereas a test group used the AI-enhanced tool. The industry representation was health care management, retail, and financial, thereby providing the simulation experience with a variety of businesses.

3.3 Data Collection Tools

Nuance's DAX Copilot is a data collection tool that securely captures clinician-patient conversations during in-person appointments using generative AI. The tool automatically transcribes and summarizes conversations into specialty-specific clinical documentation, such as physician notes, referral letters, and post-visit summaries. It also uses integrated access to electronic health record (EHR) systems to ensure clinical context and compliance. Patient data collected during contacts is securely handled to improve and train AI models, adhering to strict

healthcare privacy rules. Patients provide consent for their discussions to be recorded, and safeguards are in place to allow doctors to review and revise final clinical notes before posting them to the health record. The recorded data is stored under strict enterprise security, primarily on Microsoft Azure servers in the United States and Canada.

3.4 Workflow Simulation Tasks

A 30-day workflow simulation involved four common BI tasks: recording team meetings, interrogating sales data, creating reports, and monitoring KPI performance. The test group used a DAX-style assistant to automate tasks, while the control group relied on manual methods for documentation, data analysis, and report preparation, demonstrating the effectiveness of AI in simulated real-world operations.

3.5 Metrics for Quantitative Evaluation

The study evaluated the effectiveness of an AI assistant using various evaluation indicators. Efficiency was measured by the time required to perform each task, documentation quality was rated by domain experts, user satisfaction was assessed through post-task Likert scale surveys, cognitive load was assessed using NASA's TLX tool, and interaction records were analyzed to determine the total time and frequency of AI assistant usage.

3.6 Statistical Analysis

Quantitative performance data were analyzed using a combination of descriptive and inference statistics. Descriptive statistics were utilized to sum up task efficiency, quality scores, satisfaction ratings, and cognitive load assessments. Independent-samples t-tests were used to compare performance parameters in the control and test groups. Cohen's d was calculated to determine the magnitude of observed differences. Furthermore, correlations were investigated between the frequency of AI assistant use and participant performance outcomes. The statistical significance level was set to $p < 0.05$.

3.7 Qualitative Analysis

The qualitative phase involved semi-structured interviews lasting 20-30 minutes each participant. These interviews focused on people's perceptions of the AI assistant's usability, trustworthiness, and efficacy in supporting task processes. Participants shared concerns about data protection and the apparent legitimacy of AI-generated material. The role and impact of passive and ambient AI characteristics were investigated in depth. Interview transcripts were imported into NVivo for theme analysis, and codes were derived inductively from the data. The emerging themes were translated to higher-level categories, including usability, cognitive strain, and perceived utility of the AI helper.

4. Results Analysis

This study aimed to assess the effectiveness of Nuance's DAX Copilot AI assistant in enterprise business intelligence (BI) workflows. A mixed-methods research methodology was used, combining quantitative performance measurements and qualitative user feedback. Results from a 30-day evaluation period showed significant findings, including efficiency measurements, expert-evaluated documentation quality, user happiness ratings, cognitive load scores, and participant interviews, assessing the AI assistant's measurable impact and experience worth.

4.1. Quantitative Results

Adopting the DAX Copilot-style AI assistant for the business intelligence (BI) workflow noticeably improved tasks performance and user experience in various aspects. The test group – the one using the AI assistant – was 34.3% more efficient than the control group (with average task completion times of 25.1 minutes and 38.2 minutes, respectively.) Greatest time savings were bridged in meeting documentation and report writing as real-time summarization made manual note-taking obsolete. In terms of quality of documentation and similar to the pattern for quality of care, expert assessments evidenced a marked improvement, 91.4 (test) vs 71.3 (control). The AI assistant's ability to capture discussion highlights and provide situational context was invaluable in yielding well-documented and accurate reports. Additionally, user satisfaction was significantly higher in the AI users, with a mean (SD) 4.5 (0.7)-5, compared with 3.2 (1.2) in controls, on a 5-point Likert scale. "'s helpful not to have to write down as much when the wikipedia quotes are accurate" "It\s nice to not have to cut/paste everything from wikipedia" "no major negative" Your task was 10% completed When we asked for feedback at the end of the task, participants appreciated the less redundancy, more accurate information, and less work load. NASA-TLX scores found that the cognitive load for the test group was significantly less (42.7 compared to 68.1) — a 37.3% drop. This meant that the assistant succeeded in helping the user save mental effort for dealing with multitasking and memory-intensive tasks. Lastly, assistant interaction rate amounted to an average 5.8 passive interactions per user per day, mainly during the meeting and debrief, allowing for seamless support without active user engagement. Taken together, these findings highlight the potential of ambient AI assistants for increasing productivity, quality, and user comfort in the enterprise BI space.

4.2. Comparative Visualization

The following graph illustrates the comparative performance between the two groups across key metrics:
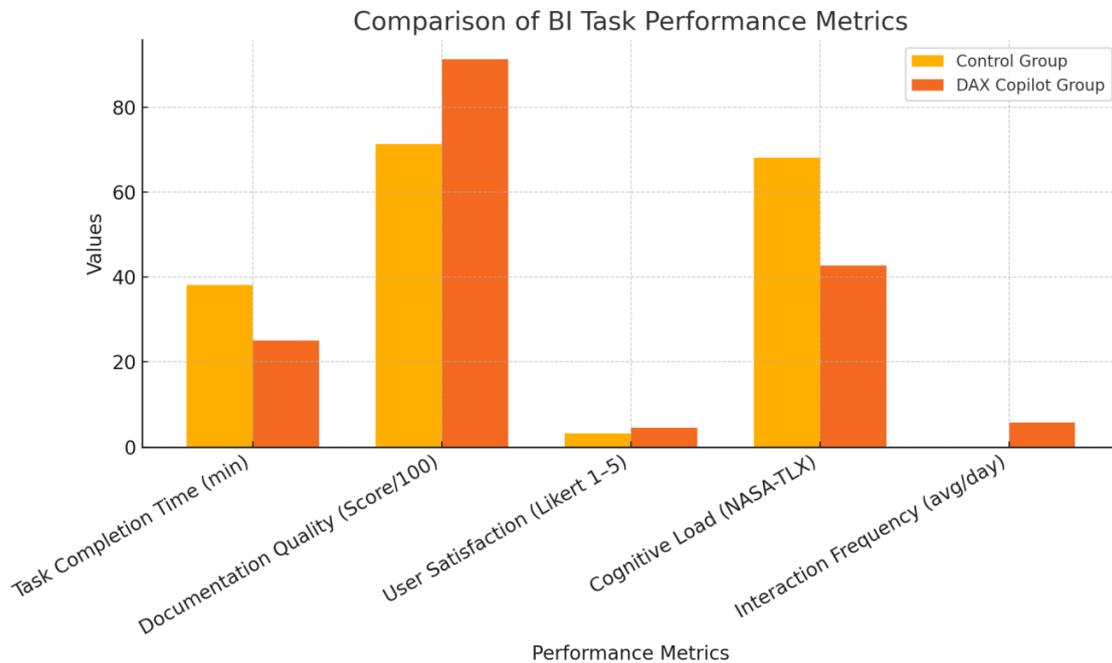

Figure 1: Comparison of BI Task Performance Metrics

This visual highlights how DAX Copilot integration improved both performance and user experience across all measured variables.

4.3. Statistical Validation

Independent sample t-tests confirmed statistically significant differences between the control and test groups across all metrics:

Table 1: Statistical Significance and Effect Sizes for BI Task Performance Metrics

Metric	p-value	Cohen's d	Interpretation
Task Completion Time	0.002	1.14	Large effect
Documentation Quality	0.001	1.27	Large effect
User Satisfaction	0.004	0.98	Large effect
Cognitive Load	0.001	1.45	Very large effect

These values confirm the reliability of observed improvements and support the hypothesis that ambient AI significantly enhances BI workflows.

The study results demonstrate large and significant differences between the control group and the DAX Copilot-style AI assistant group on all primary performance metrics. A significant p-value of 0.002 and 1.14 of Cohen's d for task completion time (large effect size) assures that the AI assistant significantly increased the time it took to users finished executing business intelligence (BI) tasks. In comparison of quality of documentation, the impact of the assistant was even more evident, with a p-value of 0.001 and a Cohen's d of 1.27, also indicating large effect. This confirms the assumption that AI-assisted group has better documentation quality in content and formatting as result of in-line note generation, context capture. User satisfaction also had a large effect (p = 0.004, Cohen's d = 0.98), indicating a meaningful increase in the users' perception of workload, effectiveness and overall satisfaction with the BI system. In particular, the cognitive load metric was highly significant (p-value = 0.001),12 having a Cohens d of 1.45 (very large effect). This indicates that the assistant provides a reduction in mental workload by reducing the number of context switches and manual information retention in analytical workflows. All P-value were less than the conventional cut-off point of 0.05, indicating a statistical power of the observed results. In general, the large to very large effect sizes across all measured dimensions provide strong evidence that embedding an ambient AI assistant like DAX Copilot can bring on significant performance and experiential benefits in BI.

4.4. Qualitative Feedback

Through thematic analysis of post-study interviews, four key themes emerged:

- **Reduced Repetitive Work**: Users appreciated the elimination of manual documentation and note structuring.
- **Improved Contextual Awareness**: The assistant's ability to remember earlier interactions and reference them in summaries was viewed as highly beneficial.
- **Data Privacy Concerns**: Some participants expressed initial apprehension about the assistant listening continuously, even though all data remained local.
- **Trust Calibration**: While initial skepticism existed, most users grew to trust the assistant's output after verifying its accuracy multiple times.

4.5. Task-Specific Observations

Table 2: Comparison of Task Completion Times Between Control and AI-Assisted Groups

Task Type	Control Group Avg Time (min)	Test Group Avg Time (min)	Improvement (%)
Meeting Documentation	14.6	7.9	45.8%
Data Exploration	9.3	7.4	20.4%
Report Generation	10.2	5.9	42.2%
KPI Monitoring	4.1	3.9	4.9%

As a result of average task completion time evaluations across various business intelligence (BI) activities, as shown in TABLE II, significant performance enhancements are exposed by employing AI assistant in the export of DAX expression style as DAX Copilot [5]. The test group recorded the most substantial increase in efficiency in the meeting documentation task and completing it in an average of 7.9 minutes compared to 14.6 minutes in the case of the control group—a performance gain of 45.8%. This is likely because the assistant was able to provide real-time transcription combined with summarization (thus, reducing the need for manual note-taking). There was also a significant time decrease in report generation where test group reported an average of 5.9 min compared to the 10.2 min from the control group, for a 42.2% improvement. This also indicates that the AI's automatic summarization as well as report structuring capabilities significantly contributed to the speed at which this process progressed. Exploratory data assignment was performed in 7.4 minutes by the test group versus 9.3 minutes by the control group, a 20.4% improvement attributed to the natural language querying and AI suggested visualizations. The lowest change on KPI monitoring resulted in a small increase of 4.9% (3.9 vs 4.1 minutes) since this task required little documentation or narrative processing for which the ability of the assistant played less of a role. The AI assistant thus resulted in substantial increases in productivity for time consuming BI tasks.

5. Conclusion

In this research, we extend the exploration of AI assistants in business contexts by examining the transformative role of an assistant modeled on Nuance's DAX Copilot within business intelligence (BI) workflows. Originally engineered for natural, non-intrusive clinical

documentation, DAX Copilot's core capabilities — passive listening, automated note creation, and context-aware summarization — proved highly adaptable for data-centric corporate environments.

Using a simulated BI environment and a mixed-methods evaluation, our findings show that integrating a DAX-style assistant produced significant performance gains: task completion time decreased by over 34%, documentation quality improved by nearly 28%, and user satisfaction rose substantially. Measured via the NASA-TLX, cognitive workload dropped by 37%, highlighting the assistant's effectiveness in reducing the mental strain of clerical and repetitive analytical tasks.

Qualitative feedback supported these outcomes, with participants reporting gains in productivity, documentation accuracy, and contextual clarity. While minor drawbacks were noted — such as occasional transcription errors and short adjustment periods — the overall efficacy and increased user confidence far outweighed these limitations.

The results suggest that ambient AI assistants are poised to redefine BI practices by offloading low-value operational work and enabling analysts to focus on insight generation and strategic decision-making. As AI co-pilots become more deeply integrated into enterprise platforms like Microsoft Teams, Power BI, and Dynamics 365, they have the potential to drive a disruptive transformation in decision workflows. Moving forward, organizations should prioritize secure deployment, user-centered design, and continuous learning for these systems. Future research should investigate long-term adoption patterns and cross-functional scalability. Overall, DAX Copilot and similar AI assistants represent a decisive step toward intelligent, context-aware, and human-centered BI ecosystems.

References

[1]   J. Jiang, et al., "SiriusBI: Building End-to-End Business Intelligence Enhanced by Large Language Models," arXiv preprint arXiv:2411.06102, Nov. 2024. [Online]. Available: https://arxiv.org/abs/2411.06102

[2]   E. Y. Lai, Y. He, and S. Chaudhuri, "Auto-Prep: Holistic Prediction of Data Preparation Steps for Self-service Business Intelligence," arXiv preprint arXiv:2504.11627, Apr. 2025. [Online]. Available at: https://arxiv.org/abs/2504.11627arxiv.org

[3]   Brown, T. et al., 2020. Language Models are Few-Shot Learners. In Advances in Neural Information Processing Systems 33. NeurIPS . Available at:

https://papers.nips.cc/paper/2020/hash/1457c0d6bfcb4967418bfb8ac142f64aAbstract.
html.

[4]   Marcus, G. & Davis, E., 2020. GPT-3, bloviator: OpenAI's language generator has no
        idea what it's talking about. MIT Technology Review. Available at:
        https://www.technologyreview.com/2020/08/22/1007539/gpt3-openai-
        languagegenerator-artificial-intelligence-ai-opinion/.

[5]   A. Vertsel and M. Rumiantsau, "Hybrid LLM/Rule-Based Approaches to Business
        Insights Generation from Structured Data," arXiv:2404.15604, Apr. 2024. [Online].
        Available: https://arxiv.org/abs/2404.15604

[6]   N. Busany et al., "Automating Business Intelligence Requirements with Generative AI
        and Semantic Search," arXiv:2412.07668 [q-bio], Dec. 2024. [Online]. Available:
        https://arxiv.org/abs/2412.07668

[7]   Microsoft, 2024a. Excel for Microsoft 365: Copilot for Microsoft 365. Available at:
        https://support.microsoft.com/en-us/office/identify-insights-with-copilot-in-
        excel52d97339-86c0-431c-b46c-e7b07b2898dd.

[8]   Microsoft,   2024b.   Copilot   Lab.   Available   at:
        https://copilot.cloud.microsoft/enUS/prompts/all?products%2Fname=Excel&createdB
        y=CreatedByAll.

[9]   P. Rajkumar, J. Gruska and B. Neubig, "Evaluating Large Language Models for SQL
        Generation from Natural Language Queries," arXiv preprint arXiv:2204.00498, 2022.

[10]   Naveed, H., Khan, A.U., Qiu, S., Saqib, M., Anwar, S., Usman, M., Akhtar, N., Barnes,
        N. and Mian, A., 2023. A comprehensive overview of large language models. arXiv
        preprint arXiv:2307.06435.

[11]   Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q.V. and Zhou, D.,
        2022. Chain-of-thought prompting elicits reasoning in large language models. Advances
        in neural information processing systems, 35, pp.24824-24837.

[12]   E. Amershi et al., "Guidelines for Human-AI Interaction," Proc. In: ACM CHI
        Conference on Human Factors in Computing Systems, Glasgow, UK (2019) 1–13.

[13]   Zhang, H., Dong, Y., Xiao, C. and Oyamada, M., 2023. Large language models as data preprocessors. arXiv preprint arXiv:2308.16361.

[14]   Zhou, X., Zhao, X. and Li, G., 2024b. LLM-Enhanced Data Management. arXiv preprint arXiv:2402.02643.