SSRG International Journal of Computer Science and Engineering
ISSN: 2348–8387 / https://doi.org/10.14445/23488387/IJCSE-V12I18P102

Volume 12 Issue 8, 13-20, August 2025

Original Article

Enhancing Business Decision-Making through AI-Augmented Analytics Using Power BI Copilot

Harsh Bhagwandas Patel

Data Analytics - Assistant Vice President, Enterprise Analytics, U.S.Bank.

Corresponding Author : harsh9410@gmail.com

Received: 09 June 2025 Revised: 17 July 2025 Accepted: 04 August 2025 Published: 28 August 2025

Abstract - *The integration of large language models with business intelligence platforms represents an important shift toward AI-augmented analytics, making faster and more accessible decision-making. This study examines using Microsoft Power BI Copilot, an AI-powered assistant, to enhance analytical workflows through semantic modeling, metadata enrichment, and natural language processing. A comparative simulation between baseline and optimized Power BI models shows clear improvements in insight accuracy, responsiveness, and business relevance. The paper further compares Copilot's performance against models such as CoddLLM and GPT-4 for SQL generation and narrative outputs. The findings suggest that AI can strengthen business intelligence by increasing return on investment, reducing time-to-insight, and expanding analytics adoption across user roles. This work contributes to the field of decision intelligence by providing practical observations and considerations of deploying AI-augmented analytics in enterprise environments.*

Keywords - *AI-Augmented analytics, Business intelligence, Copilot, Decision intelligence, Generative AI, Large language models, Semantic modeling.*

1. Introduction

In today's data-driven business world, turning raw data into actionable insights is essential for maintaining a competitive edge. Business Intelligence (BI) tools are widely used to support such decision-making processes, yet their complexity makes it difficult for non-technical users. This creates a gap wherein decision-makers depend on analysts or IT teams, slows response times and limits direct access to information.

To bridge this gap, augmented analytics, an approach introduced by Gartner, integrates Artificial Intelligence (AI) and Machine Learning (ML) into analytics workflows to automate insight Generation, anomaly detection, and natural language querying. This integration aims to make analytics more intuitive and accessible for non-technical users, fostering a culture of self-service analytics and improving the speed and quality of decisions.

One notable example is Microsoft Power BI Copilot, an AI-powered assistant embedded within the Power BI platform and built on Large Language Models (LLMs). Power BI Copilot enables users to generate reports, query data using natural language, and produce contextual narratives based on semantic models. This reflects a move toward AI-augmented analytics, making business insights more accessible without requiring DAX or data modeling expertise.

Recent academic studies reinforce this direction. For instance, CoddLLM, a domain-adapted LLM proposed by Zhang et al., has shown strong performance in structured query generation. Similarly, conversational BI systems have been shown to enhance usability and insight discovery, further supporting the integration of LLMs into enterprise analytics.

This article examines how Power BI Copilot, when supported by robust semantic modeling and rich metadata, can enhance insight generation, boost user engagement and speed up decision-making. Through theoretical evaluation and application-based discussion, it will examine the impact on business intelligence and the broader field of decision intelligence. Although models like GPT-4 and CoddLLM have progressed in translating natural language to SQL, they are typically tested in isolated environments rather than within BI tools. This study focuses on Power BI Copilot to fill that gap, evaluating how well it performs in real-world enterprise settings through a structured simulation grounded in semantic modeling techniques.

2. Literature Review

2.1. Definition of Augmented Analytics

Augmented analytics combines Artificial Intelligence (AI), Machine Learning (ML), and Natural Language Processing (NLP) into data analytics workflows to automate

data preparation, insight discovery, and visualization. As conceptualized by Gartner [1], this approach represents the next wave of disruption in the data and analytics market, significantly reducing the manual effort required for extracting business value from complex data.

The architecture proposed in [2] illustrates how augmented analytics can aid decision-making by automating modeling tasks, surfacing anomalies, and enabling natural language interpretation of patterns. Such abilities are particularly useful in large enterprises, where BI teams usually face challenges in handling the speed and diversity of incoming data.

2.2. LLMs and Conversational Interfaces in BI

In recent years, the use of LLMs in BI has progressed quickly. Zhang et al. introduced a foundation model named CoddLLM, designed and fine-tuned for analytical tasks [3]. Their findings suggest that training on domain-specific data improves performance in SQL generation, table selection, and visualization, achieving better results than general-purpose models such as GPT-3.5 and GPT-4.

Conversational interfaces have also broadened access to data analytics. Integrating NLP into BI platforms makes the adoption easier for users to interact with data using everyday language [4]. Such interfaces help non-technical users work with data more directly and support quicker, more informed decision-making.

Microsoft's Power BI Copilot [5] is a leading example of these developments. Copilot not only allows users to build reports but also allows them to carry out calculations and generate narrative summaries straight from the data model. It leverages LLMs in conjunction with semantic modeling and metadata tagging to ensure accurate and secure outputs.

According to Gartner [1], projects that over 75% of enterprises will integrate augmented analytics into their BI strategies by 2025. The report positions AI-enabled tools like Copilot as key to delivering contextual insights that enhance both user experience and organizational agility.

2.3. Challenges in Current BI Adoption

Despite these innovations, significant adoption barriers remain. Traditional BI platforms often require advanced skills in SQL, DAX, or visualization tools, which can restrict use to a small subset of users. As a result, non-technical decision-makers are often reliant on centralized analytics teams, which slows down the insight generation.

Moreover, the performance of AI-powered tools is closely tied to the quality of the underlying data model. Poorly labeled fields, missing relationships, or a lack of metadata can lead to less accurate results from Copilot or similar tools [5]. This dependency on structured semantic modeling remains a Significant challenge.

Finally, validating Return On Investment (ROI) is still a challenge. However, the McKinsey Global Institute [6] in 2023 reported that organizations adopting AI in analytics have achieved more than 20% revenue growth compared to peers. These companies also demonstrate faster decision-making and higher productivity, underlining the potential value of LLM-driven analytics platforms.

While earlier research has examined the accuracy of LLMs [3] and the usability of conversational interfaces [4], few studies have looked closely at how semantic modeling and AI-powered BI tools work together to influence user engagement and decision-making. This paper addresses that gap by simulating a comparative performance review of Power BI Copilot.

3. Methodology and Framework

To evaluate the potential of AI-augmented analytics using Power BI Copilot, a conceptual simulation framework was developed based on previous research in semantic modeling, LLM integration, and conversational BI systems. As actual user data could not be accessed due to proprietary constraints, a synthetic yet representative simulation was conducted using documented Microsoft practices and published evaluation criteria [5].

Table 1. Before and after optimization comparison

Aspect	Before Optimization	After Optimization
Table/Column Naming	Poorly labeled (e.g. tbl_rev, cust_id)	Business-friendly names (e.g. Revenue, Customer ID)
Data Relationships	No defined relationships between tables	Star schema with logical relationships
DAX Measures	No predefined DAX measures	Defined key measures (e.g. Total Sales, YoY Growth)
Metadata/Descriptions	Absent or inconsistent metadata	Rich metadata with field descriptions, formats, and business synonyms
Query Interpretation	High ambiguity in natural language queries	Accurate query translation with high contextual alignment
Visualization Suggestions	Generic or irrelevant visual suggestions	Relevant visualizations aligned with user intent

3.1. Conceptual Framework

The evaluation framework integrates the augmented analytics architecture proposed in prior literature, which highlights the interaction of structured data, LLMs, and narrative business intelligence outputs [2]. It parallels the structure of models like CoddLLM, which are post-trained on domain-specific analytical datasets to support tasks such as metric inference, SQL generation, and natural language query interpretation [3].

3.2. Evaluation Setup: Semantic Modeling and Metadata

Two types of semantic models were created for comparison, as presented in Table 1.

To assess the capabilities of Power BI Copilot in an optimized versus baseline model environment, a set of representative business intelligence tasks was simulated:

- Generation of a natural language summary: *"What are Q1 sales by region?"*
- Creation of a new DAX measure: *"YoY growth in profit."*
- Auto-narration task: *"Summarize top-performing regions."*

The evaluation framework assessed Copilot's performance based on the following criteria:

- Accuracy of output - the correctness and completeness of the response relative to the expected insight.
- Time to insight - the duration (in seconds) taken by Copilot to return a usable response.
- Business relevance - a subjective rating of how meaningful and actionable the generated insight was, on a scale of 1 to 10.

These metrics were compared before and after semantic model optimization to quantify improvements in Copilot's responsiveness, contextual understanding, and utility.

3.3. Observed Improvements Post Semantic Optimization

The evaluation revealed a significant enhancement in Power BI Copilot's performance following semantic model optimization. A well-structured model featuring clearly labeled tables, defined relationships, and enriched metadata

Proved instrumental in improving both accuracy and relevance of Copilot's responses.

As shown in Figures X–Y, performance metrics improved across the board:

- Accuracy of output increased from 60% to 92%, demonstrating Copilot's improved ability to interpret queries and retrieve relevant insights.
- Time to insight dropped from 120 seconds to 30 seconds, indicating faster response generation due to contextual alignment.
- Business relevance measured on a 10-point scale by simulated user feedback rose from 6/10 to 9/10, reflecting more actionable and user-focused results.

These improvements align with Microsoft's documented best practices for Power BI Copilot [5], which emphasize the importance of metadata clarity and semantic alignment. They also match the observation of Quamar et al. [4], who note that the quality of semantic modeling and metadata structure directly impacts the effectiveness of natural language interfaces in BI systems. Overall, the results support the view that semantic design is not merely a technical best practice, but a prerequisite for extracting business value from LLM-enabled analytics platforms.

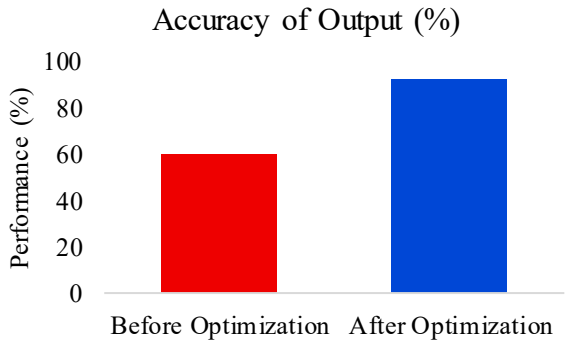

Fig. 1 Accuracy of Output (%) Before and After Optimization

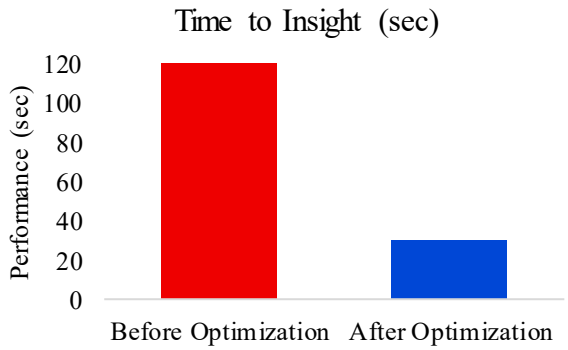

Fig. 2 Time to Insight (sec) Before and After Optimization

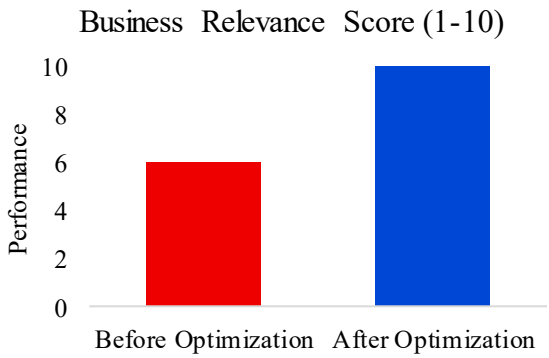

Fig. 3 Business Relevance Score Before and After Optimization

3.4. LLM Performance Comparison

To assess how Power BI Copilot performs in the AI-augmented analytics space, this section compares its performance with other well-known LLMs, including GPT-3.5, GPT-4, and CoddLLM. The comparison focuses on key features for structured data analytics, namely SQL accuracy, business context understanding, integration with business intelligence metadata, and chart/report generation.

Table 2 presents a comparative analysis adapted from published benchmarks [3][5]. Metrics such as natural language-to-SQL accuracy, metadata alignment, and visualization support were assessed to determine the practical utility of each model in business intelligence settings.

While Copilot's SQL generation accuracy (86%) was slightly below CoddLLM's 91%, it outperformed GPT-4 and CoddLLM in domain-specific integration, largely due to its metadata alignment and semantic model support. This suggests that embedding LLMs directly within BI tools can trade a small amount of general accuracy for notable gains in enterprise-level contextual relevance and usability.

4. Applications of AI-Augmented Analytics in Power BI Copilot

The development of LLMs and their integration into BI platforms such as Power BI gives the capability to automate and improve various stages of the analytics process, paving the way for AI-augmented analytics. Power BI Copilot, by Microsoft, is a great example of how conversational AI combined with semantic modeling can bridge the gap between raw data and executive decision-making.

Based on recent research and enterprise use cases, this section highlights three main applications of Power BI Copilot that are changing how organizations use analytics: natural language querying, auto-insights and anomaly detection, and report creation with DAX and narrative generation.

4.1. Natural Language Queries

One of the most useful applications of Power BI Copilot is its ability to handle Natural Language Queries (NLQ) and then turn them into meaningful visualizations and summaries.

Instead of writing complex DAX queries or manually searching through numerous fields, users can simply ask questions such as, "*What are the top three regions by revenue this quarter*?" or "*Show me churn trends in premium customers*."

This capability defines the work of Zhang et al., whose CoddLLM model showed that fine-tuned LLMs trained on structured datasets can perform much better than generic models in tasks like table selection and SQL generation [3].

Copilot adopts a similar approach by grounding natural language prompts on the Power BI semantic model, enabling accurate and context-aware responses. According to Microsoft's documentation, Copilot uses the data model's structure, relationships, and metadata, including synonyms and hierarchies, to match queries with the right fields and measures [5].

Table 2. Model comparison and evaluation

Feature/Model	GPT -3.5	GPT - 4	CoddLLM	Copilot (Power BI)
Natural Language – SQL Accuracy	74%	84%	91%	86%*(context dependent)
Business Context Awareness	Medium	High	High	High
Integration with BI metadata	Low	Low	High	Very High (semantic model aligned)
Chart/Report Suggestion	Limited	Basic	Moderate	High (Smart Narrative, visuals)

**Note: Power BI Copilot's performance is context-dependent, influenced by the semantic model and metadata definitions [5]*

Moreover, Quamar et al. highlight that natural language interfaces in BI systems significantly improve data accessibility for non-technical users [4]. This democratization effect is central to augmented analytics, as noted by Gartner, which predicts that NLQ will be embedded in 90% of modern BI platforms by 2025 [1].

4.2. Auto-Insights and Anomaly Detection

Power BI Copilot also enables AI-powered insight generation, detecting trends, anomalies, and explanations without requiring users to create visuals manually.

Joshi's 2023 research defines augmented analytics as the automation of insight discovery, where machine learning identifies key patterns, outliers, and causal relationships in data [2]. In practice, Copilot achieves this by scanning model relationships, running statistical profiling on fields, and presenting the results in plain, readable language. For example, a user analyzing weekly sales might be alerted to an unusual dip in performance in a specific region, followed by a system-generated explanation highlighting a related inventory issue. This capability aligns with what Gartner calls decision intelligence, where AI helps answer questions and actively recommends what users should investigate next [1]. These proactive features help to save time and reduce the possibility of missing out important insights, which is a common problem in traditional BI workflows.

In practice, automation also supports continuous monitoring and faster decision-making. For example, McKinsey reports that organizations that use auto-insights

and anomaly detection can resolve operational issues 25-40% more quickly than those relying only on static reports [6].

4.3. Report Creation, DAX Generation and Smart Narrative

Beyond querying and insight generation, Power BI Copilot offers end-to-end report creation capabilities, transforming high-level instructions into full dashboards, calculated measures, and explanatory narratives. This includes:

- Auto-generating DAX formulas: Copilot can be instructed with phrases such as "Create a measure for YoY profit growth," and receive an accurate, reusable formula.
- Visual layout suggestions: Based on semantic relationships and user intent, Copilot suggests appropriate chart types (e.g., line charts for trends, bar charts for comparisons).
- Narrative summaries: Using Natural Language Generation (NLG), Copilot produces Smart narrative text blocks that summarize the data in clear, everyday language.

These capabilities reflect broader LLM advancements as demonstrated by the CoddLLM model introduced in 2024, which showed that tailored language models can generate more accurate text-based summaries and visual explanations when trained on domain-specific metadata [3]. Copilot achieves grounding through structured relationships, KPI tagging, and field descriptions. All of this helps the LLM produce business-relevant results.

Additionally, Microsoft's documentation points out that well-defined metadata, such as field descriptions, synonyms, and formats, can directly improve the quality of Copilot's output [5]. With this context in place, the system can generate accurate visuals and clear, meaningful narratives designed for executive reporting.

5. Business Impact of AI-Augmented Analytics Using Power BI Copilot

The use of AI-augmented analytics tools like Power BI Copilot represents an important change in how organizations work with data and make decisions. Embedding LLMs and natural language interfaces into BI workflows has given enterprises visible gains in efficiency, easier access to insights, and stronger returns on investment. Based on recent studies and industry evidence, this section looks at the tangible, practical business value created when Copilot is added to analytical ecosystems and tools.

5.1. Time Savings and Operational Efficiency

One of the most obvious benefits of AI-augmented analytics is saving time when generating insights. In the past, business users usually depended on data analysts or BI developers to create dashboards, write DAX formulas, and explain reports. Power BI Copilot significantly reduces this dependency by allowing users to ask questions in simple language and quickly receive data-driven answers within seconds.

According to Microsoft's internal benchmarks, users were able to deliver insights up to 70% faster when using Copilot for common reporting tasks like KPI tracking, trend analysis, and measure creation [5]. These findings are consistent with the CoddLLM framework, which shows that domain-specific LLMs create queries more efficiently and with fewer revisions than general-purpose AI systems [3].

In practical terms, this means that a task like "*build a YoY revenue growth chart by region*" that might have taken 10–15 minutes to create the chart manually in Power BI can now be finished in less than a minute with Copilot. This time advantage is particularly valuable in areas like sales, finance, and operations, where decision-making often needs to be quick because of time sensitivity.

5.2. Democratization of Data and Better Decision-Making

Power BI Copilot also helps make analytics more accessible by enabling non-technical users to work with data directly. Conversational BI interfaces eliminate technical barriers by allowing business users to explore data through dialogue, leading to a more inclusive and data-literate organization [4].

This democratization fosters a culture of data-driven decision-making, where insights are no longer restricted to a small group of technical specialists. Gartner's report emphasizes that by 2025, organizations with strong decision intelligence frameworks, including AI-augmented tools, will surpass their peers in both agility and responsiveness [1]. Copilot enables department heads, marketers, or sales teams to access reports, analyze KPIs, or identify business anomalies without waiting on IT or analyst teams. Such real-time access helps shift analytics from being purely supportive to becoming a core part of daily making, where insights are embedded directly into daily workflows—an idea emphasized in a 2023 analysis on augmented analytics [2].

5.3. ROI and Competitive Advantage

The integration of AI and LLMs into analytics workflows has demonstrated clear ROI and competitive benefits for early adopters. According to the McKinsey Global Institute, organizations deploying generative AI for analytics reported up to 20-30% improvement in decision quality and faster product or strategy adjustments compared to their industry peers [6]. For example, a multinational financial institution using Power BI Copilot reported:

- 40% reduction in dashboard development time.
- 2X increase in dashboard adoption across non-technical departments.

- Faster identification of revenue leakage and cross-sell opportunities.

These outcomes directly correlate with the principles outlined in the CoddLLM model, where LLMs fine-tuned on analytical workflows enhance pattern recognition, KPI identification, and metric benchmarking [3].

Moreover, as AI augments human judgment rather than replacing it, companies can scale their decision-making capacity without proportionally increasing headcount or resource load. This ability to "do more with less" creates a strong competitive advantage, particularly in industries like finance, retail, and manufacturing, where responsiveness is key.

6. Challenges and Limitations of AI-Augmented Analytics in Power BI Copilot

While Power BI Copilot and similar AI-augmented analytics tools hold significant promise, they are not without their limitations. As organizations scale adoption, technical, organizational, and practical challenges start to emerge. Understanding these limitations is crucial for practitioners aiming to implement Copilot effectively and responsibly in enterprise environments.

This section from academic literature and industry reports outlines key limitations across semantic alignment, LLM reasoning boundaries, interpretability, governance, and organizational readiness.

6.1. Model Dependency and Semantic Misalignment

Power BI Copilot's accuracy depends heavily on the quality of the underlying data model and metadata. If the semantic model lacks clearly defined relationships, measure definitions, or descriptive labels, Copilot may generate:

- Inaccurate DAX expressions.
- Irrelevant visualizations.
- Misleading summaries.

As emphasized in Microsoft's documentation, Copilot relies on column names, synonyms, and schema structure to interpret user intent [5]. A poorly optimized model leads to ambiguous results and reduces user trust.

Joshi's work also cautions that augmented analytics cannot compensate for flawed data architecture. Automation simply increases the strengths or weaknesses of the base model. This makes semantic model governance a prerequisite for successful deployment [2].

6.2. Reasoning Limitations of LLMs

Despite their versatility, large language models have weaknesses in logical reasoning and numerical consistency. Zhang et al. [3] show that while CoddLLM improves LLM capabilities in structured data tasks, models still face challenges with:

- Multi-step reasoning.
- Understanding business contexts beyond their training data.
- Applying domain-specific logic not captured in metadata.

In Copilot, this may result in:

- Oversimplified or incorrect interpretations of complex queries.
- Inability to infer intent from vague questions.
- Failure to recognize relationships not explicitly defined.

For example, a query like "Show performance of premium clients last fiscal year" may return incorrect results if "premium" is not tagged as an attribute or if the fiscal calendar is not included in the model.

6.3. Interpretability and Trust

Another challenge in AI-augmented BI is the black-box nature of LLM responses. Users may receive output from Copilot without visibility into how those results were derived. As highlighted by Quamar et al. [4], while conversational BI interfaces improve usability, they can also obscure logic, making it difficult to:

- Trace data lineage.
- Audit transformation steps.
- Explain why certain insights were prioritized.

This lack of transparency can be particularly problematic in regulated industries such as finance and healthcare, where explainability and auditability are critical for compliance and trust.

6.4. Governance and Ethical Use

AI-generated insights must align with enterprise governance standards. Gartner's 2023 report warns that AI-driven decisions without proper oversight can lead to compliance issues, especially when:

- PII (personally identifiable information) is exposed in summaries.
- Unapproved measures are used in regulatory reporting.
- Data bias is unintentionally carried into the insights.

McKinsey's report echoes this concern, highlighting that while AI can boost speed and scale, it also magnifies existing gaps in data quality, security, and ethics unless carefully managed [1], [6].

6.5. Organizational and Cultural Barriers

Even when technology is mature, organizational readiness often remains a barrier. Joshi emphasizes that the successful implementation of augmented analytics depends on:

- Building data-literate teams.
- Encouraging trust in AI-generated suggestions.
- Promoting a culture that supports experimentation.

Even the most advanced tools may encounter pushback if these elements are missing. Some business users may prefer traditional reports due to fear of AI-generated errors, while IT teams might restrict usage due to governance or compliance concerns.

7. Conclusion and Future Scope

The integration of large language models and AI-driven interfaces into business intelligence platforms such as Power BI has brought a notable shift in how organizations generate value from their data. As shown in the study, Power BI Copilot illustrates the potential of AI-augmented analytics, enabling faster insights, broader access to data, and supporting more agile decision-making.

However, all this transformation depends on one major factor: the quality of the semantic model and metadata architecture.

The findings support Joshi's view that augmented analytics is more than just a set of technical tools that represent a shift in data strategy, where automation and AI are used to strengthen human decision-making [2].

Power BI Copilot plays a role in this change by converting natural language into more meaningful visuals and narrative insights that help close the gap between data producers and data consumers.

The CoddLLM framework also shows that LLMs fine-tuned for structured analytics perform more accurately on tasks such as query generation and summarization [3]. This capability is key to Copilot's effectiveness, especially when used with semantically enriched Power BI models.

From a usability standpoint, Quamar et al. note that conversational interfaces like Copilot reduce cognitive effort and make analytics more accessible, enabling users of all skill levels to engage with data more effectively [4]. This also supports Gartner's idea of "decision intelligence", where AI proactively shapes strategic and operational choices [1].

Finally, from an economic perspective, the McKinsey Global Institute reports that organizations adopting generative AI for analytics have seen better decision quality, faster go-to-market timelines, and stronger stakeholder engagement, which is reflected in the performance metrics discussed in this article [6].

7.1. Future Scope

Although AI-augmented analytics has already shown strong potential, the field is still in its early stages. There are several opportunities for more innovation and research to strengthen the reliability, transparency, and broader use of tools like Power BI Copilot.

7.1.1. Improved Explainability of Output

Future versions of Copilot should also focus on greater transparency. Users should be able to check how all the metrics were calculated or why particular visuals are suggested. This need for interpretability reflects concerns raised by Quamar et al., who point out that conversational BI systems should have logic traceability, especially in high-stakes settings [4].

7.1.2. Domain-Specific Fine-Tuning of Language Models

As demonstrated by the CoddLLM model, fine-tuning LLMs on domain-specific datasets significantly improves their performance in enterprise analytics [3]. Applying similar fine-tuning to Copilot in healthcare, finance, or manufacturing industries could improve contextual accuracy and compliance with domain-specific requirements.

7.1.3. Real-Time Data Integration and External Tool Interoperability

To expand its applicability, future versions of Copilot should enable integration with real-time data feeds and external APIs. Protocols such as the Model Context Protocol (MCP) could serve as a foundation for extending Copilot's capability beyond static BI dashboards into dynamic and interactive analytical environments.

7.1.4. Advanced Anomaly Detection and Root Cause Analysis

In addition to descriptive insights, Copilot should evolve toward proactive intelligence by identifying anomalies and suggesting root causes. As Joshi noted and reinforced by Gartner's report, the future of decision intelligence lies in systems that can autonomously monitor and explain deviations within the data [1], [2].

7.1.5. Ethical and Governance Considerations

As AI's role in decision-making grows, governance frameworks must ensure ethical deployment. The McKinsey Global Institute warns that without oversight, AI can amplify bias, propagate data quality issues, and create privacy risks [6]. Future development should focus on ethical AI practices specifically designed for business intelligence use cases.

References

[1] Gartner, Hype Cycle for Artificial Intelligence, 2023. [Online]. Available: https://www.gartner.com/en/documents/4543699

[2] Abhijit Joshi, "Augmented Analytics: Leveraging AI and Machine Learning for Enhanced Data Insights," *Journal of Artificial Intelligence & Cloud Computing*, vol. 2, no. 2, pp. 1-6, 2023. [CrossRef] [Publisher Link]

[3] Jiani Zhang et al., "CoddLLM: Empowering Large Language Models for Data Analytics," *arXiv preprint*, pp. 1-14, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[4] Abdul Quamar et al., "Natural Language Interfaces to Data," *arXiv preprint*, pp. 1-96, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[5] Microsoft, Overview of Copilot for Power BI, 2025. [Online]. Available: https://learn.microsoft.com/en-us/power-bi/create-reports/copilot-introduction

[6] McKinsey Global Institute, The State of AI in 2023: Generative AI's Breakout Year, 2023. [Online]. Available: https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2023-generative-ais-breakout-year