Integración de Modelos de Lenguaje de Gran Escala con Power BI para la Generación Automática de Insights

Integration of Large Language Models with Power BI for Automatic Insight Generation

DOI: http://dx.doi.org/10.17981/cesta.06.01.2025.06
Scientific research article
Date received: February 25, 2025, Date accepted: April 15, 2025.

Dagoberto Altamar Pacheco
Universidad del Sinú Seccional Cartagena (Colombia)
dagobertoaltamar@unisinu.edu.co

Pedro Luis Torres Alvarez
Universidad del Sinú Seccional Cartagena (Colombia)
Pedro.torresa@unisinu.edu.co

How to cite:
D. Altamar Pacheco and P. L. Torres Alvarez, "Integration of Large Language Models with Power BI for Automatic Insight Generation", J. Comput. Electron. Sci.: Theory Appl., vol. 6, no. 1, pp. 25-33, 2025. DOI: 10.17981/cesta.06.01.2025.06

Abstract

Introduction: The growing availability of corporate data has increased the need for tools that facilitate analysis without requiring advanced technical expertise, fostering the integration of generative artificial intelligence into business intelligence platforms such as Power BI. **Objective:** To analyze the integration of large language models with Power BI to enable augmented analytics capabilities focused on automatic insight generation. **Method:** A conceptual review and analytical approach was conducted on business intelligence, the Transformer architecture, and automatic insight generation, complemented by the proposal of a five-layer architecture and the review of implementation patterns using Power BI Copilot, Azure OpenAI, and custom APIs. **Results:** Key capabilities were identified, including natural language-to-query translation, automatic DAX code generation, explanatory narrative creation, and conversational interaction with data, as well as applications in retail, human resources, and supply chain. **Conclusions:** The integration of large language models with Power BI represents a significant step toward democratizing data analysis by enabling non-specialized users to access insights in an agile, understandable, and effective manner.

Keywords

Power BI, large language models, generative artificial intelligence, augmented analytics, business intelligence, automatic insight generation.

Resumen

Introducción: La creciente disponibilidad de datos corporativos exige herramientas que faciliten su análisis sin requerir conocimientos técnicos avanzados, impulsando la integración de inteligencia artificial generativa en plataformas de inteligencia de negocios como Power BI. **Objetivo:** Analizar la integración de modelos de lenguaje de gran escala con Power BI para habilitar capacidades de analítica aumentada y generación automática de insights. **Método:** Se realizó una revisión y análisis conceptual sobre inteligencia de negocios, arquitectura Transformer y generación automática de insights, complementado con una arquitectura de cinco capas y patrones de implementación con Power BI Copilot, Azure OpenAI y APIs personalizadas. **Resultados:** Se identificaron capacidades como traducción de lenguaje natural a consultas, generación automática de código DAX, narrativas explicativas e interacción conversacional con datos, además de aplicaciones en retail, recursos humanos y cadena de suministro. **Conclusiones:** La integración de estos modelos con Power BI favorece la democratización del análisis de datos, permitiendo que usuarios no especializados accedan a insights de forma ágil, comprensible y efectiva.

Palabras clave

Power BI, modelos de lenguaje de gran escala, inteligencia artificial generativa, analítica aumentada, inteligencia de negocios, generación automática de insights.

INTRODUCTION

Business Intelligence (BI) has undergone a significant transformation over the past decades, evolving from traditional *On-Line Analytical Processing* (OLAP) systems primarily focused on multidimensional data aggregation and reporting toward modern analytical ecosystems that emphasize interactivity, scalability, and user autonomy. Contemporary BI platforms enable advanced data visualization, real-time analytics, and self-service capabilities, allowing organizations to leverage data as a strategic asset for decision-making.

Despite these advancements, a critical limitation remains: the accessibility of advanced analytical capabilities. Extracting complex insights from corporate data typically requires proficiency in declarative query languages such as DAX (*Data Analysis Expressions*) or SQL (*Structured Query Language*), as well as an understanding of data modeling concepts, including relationships, measures, and evaluation contexts. These requirements create a barrier that restricts deep analytical exploration to technically trained users, limiting the democratization of data and slowing down decision-making processes across organizational levels.

In parallel, the rapid growth in data volume, velocity, and variety has intensified the need for more intuitive and scalable analytical tools. Organizations increasingly demand solutions that not only process large datasets efficiently but also translate them into actionable insights in a timely and understandable manner. In this context, the emergence of Large Language Models (LLMs), based on architectures such as the Transformer [8], represents a disruptive advancement in artificial intelligence.

LLMs, including models such as GPT-4, Claude, and Llama, exhibit advanced capabilities in natural language understanding, contextual reasoning, and code generation [7], [15]. These capabilities allow them to interpret user intent, generate executable analytical queries, and produce coherent explanatory narratives. As a result, LLMs enable a new paradigm in which users can interact with data systems using natural language rather than formal query syntax.

In response to this technological shift, Microsoft has incorporated generative AI capabilities into Power BI through Copilot [1]. This integration embeds Natural Language Processing (NLP) directly into the analytical workflow, enabling users to formulate queries conversationally and receive automatically generated visualizations, measures, and insights. By abstracting the complexity of query languages and data modeling, this approach significantly reduces the technical barrier to data analysis.

Furthermore, this convergence between BI platforms and LLMs aligns with the broader concept of augmented analytics, where artificial intelligence automates key stages of the analytical lifecycle, including data preparation, insight generation, and result interpretation [10]. This shift not only enhances efficiency but also transforms the role of users, moving from manual data manipulation toward strategic interpretation and decision-making.

Consequently, the integration of LLMs into BI platforms represents more than a technological enhancement—it constitutes a paradigm shift in how organizations interact with data. It redefines the relationship between users and analytical systems, promotes a more inclusive data culture, and enables faster, more informed decision-making across all organizational levels.

THEORETICAL FRAMEWORK

A. Foundations of Business Intelligence and the Power BI Platform

Business Intelligence (BI) encompasses a comprehensive set of methodologies, technologies, and processes designed to transform raw data into meaningful and actionable information. At its core, BI aims to support strategic and operational decision-making by enabling organizations to analyze historical and real-time data, identify trends, and monitor performance indicators.

The evolution of BI has been closely linked to advancements in data warehousing and dimensional modeling. Foundational approaches such as those proposed by Kimball emphasize the use of star and snowflake schemas to structure data in a way that facilitates efficient querying and analysis [14]. These models organize data into fact tables and dimension tables, enabling flexible aggregation and intuitive exploration.

Within this context, Power BI has emerged as a leading platform, consistently positioned in the Magic Quadrant for Analytics and BI Platforms by Gartner [2]. Power BI integrates multiple functionalities into a unified ecosystem, including data ingestion, transformation, modeling, visualization, and sharing.

The platform is structured around three primary components:

- **Power BI Desktop**, which provides a development environment for data modeling and report creation, including ETL processes and advanced visual design.

- **Power BI Service**, a cloud-based platform that enables collaboration, report distribution, access control, and integration with enterprise systems.

- **Power BI Mobile**, which facilitates access to dashboards and reports on mobile devices, supporting decision-making in dynamic environments.

A key element of Power BI is its semantic modeling layer, which defines relationships between datasets and encapsulates business logic through calculated measures. This layer relies heavily on DAX, a functional language designed for analytical expressions [6], [11]. While DAX provides powerful capabilities for defining complex calculations, its syntax and evaluation context mechanisms particularly filter context and row context, introduce a steep learning curve for non-technical users.

Moreover, the effectiveness of BI systems depends on the quality and structure of semantic models. Well-defined metadata, consistent naming conventions, and properly designed relationships are essential for ensuring accurate and efficient query execution [5]. Poorly designed models can lead to incorrect insights, performance issues, and reduced user trust.

Therefore, while platforms like Power BI have significantly advanced the accessibility of analytics, they still require a level of technical expertise that limits their full adoption across all user profiles.

B. Large Language Models and Transformer Architecture

Large Language Models (LLMs) represent a major breakthrough in artificial intelligence, driven by the introduction of the Transformer architecture [8], which replaces traditional sequential processing with attention-based mechanisms. This architecture enables models to capture long-range dependencies and contextual relationships within text, significantly improving performance in natural language tasks.

LLMs are trained on massive datasets using self-supervised learning, allowing them to develop a deep understanding of language patterns, semantics, and structure. As demonstrated in foundational work by Brown et al. [7], these models exhibit emergent capabilities, including few-shot and zero-shot learning, where they can generalize to new tasks with minimal or no additional training.

Modern models such as GPT-4 extend these capabilities further, demonstrating strong performance in tasks that require logical reasoning, code generation, and structured problem-solving [15]. These characteristics make LLMs particularly suitable for integration into analytical systems, where they can act as intermediaries between human language and formal query languages.

In the context of BI, LLMs enable several transformative functionalities:

- They can interpret natural language queries and map them to structured analytical operations.

- They can generate executable code in languages such as SQL, Python, and DAX.

- They can analyze structured datasets and identify patterns, trends, and anomalies.

- They can produce human-readable explanations that contextualize analytical results.

These capabilities position LLMs as a bridge between technical and non-technical users, enabling a more inclusive and intuitive approach to data analysis.

C. Automatic Insight Generation

Automatic Insight Generation (AIG) refers to the use of computational methods to identify meaningful patterns and relationships in data without explicit user guidance. Traditional AIG approaches rely on statistical techniques and machine learning algorithms to detect correlations, outliers, and trends.

However, these approaches are often limited in their ability to contextualize findings or communicate them effectively to users. This is where the integration of LLMs introduces a significant advancement.

By combining statistical detection with natural language generation, LLM-enhanced AIG systems can:

- Translate quantitative findings into qualitative explanations.

- Generate hypotheses about potential causes of observed patterns.
- Adapt communication to different audiences, from technical analysts to business executives.
- Support iterative exploration through conversational interfaces.

This transformation aligns with the concept of augmented analytics, where AI technologies automate and enhance the analytical process [10]. Instead of merely presenting data, augmented analytics systems actively assist users in understanding and interpreting information.

As a result, AIG evolves from a purely analytical process into a hybrid system that integrates computation, reasoning, and communication. This shift has profound implications for organizations, as it enables faster insight generation, reduces reliance on specialized roles, and promotes a more data-driven culture.

ARCHITECTURE FOR INTEGRATING LLMS INTO BUSINESS INTELLIGENCE

This section presents a five-layer architecture and the sequential workflow that underpin Augmented Analytics through the integration of Large Language Models (LLMs) into Business Intelligence (BI) platforms, exemplified by Power BI Copilot.

A. Architectural Components of the Integration

The system operates through a set of interconnected functional layers that collectively manage natural language interaction, LLM processing, and analytical execution:

- **Presentation Layer (Front-end):** This layer provides the user interface for interacting with Copilot, including natural language query input panels and the visualization of generated outputs such as charts, dashboards, and explanatory narratives. It is designed to offer an intuitive and conversational user experience, reducing the cognitive load associated with traditional query formulation.
- **Orchestration Layer:** Acting as middleware, this layer is responsible for intent interpretation, conversational context management, and routing of requests to the appropriate LLM endpoints. It also coordinates multi-step workflows, ensuring consistency between user input, model responses, and backend execution.
- **LLM Processing Layer:** This layer hosts the core language models, typically deployed in secure environments such as Azure OpenAI Service. It incorporates prompt engineering strategies, template management, and response validation mechanisms to ensure accurate and context-aware code and text generation.
- **Data Layer:** This layer encompasses the semantic model of Power BI, including schema metadata such as tables, columns, measures, and relationships. This structured metadata is essential for grounding the LLM, enabling it to generate queries that are both syntactically correct and semantically aligned with the underlying data model.
- **Security and Governance Layer:** This layer ensures compliance with organizational and regulatory requirements by managing Row-Level Security (RLS), authentication and authorization protocols, data access policies, and content moderation. It plays a critical role in safeguarding sensitive information and maintaining trust in AI-driven analytics.

B. Data Flow and Sequential Processing

The process for automatic insight generation through Natural Language Processing (NLP) follows a structured sequence of seven stages:

1. **Natural language query capture:** The user submits a question or request in conversational language.
2. **Semantic intent analysis:** The system interprets the query, identifying user intent and extracting key entities and parameters.
3. **Contextual grounding:** Relevant metadata is retrieved from the semantic data model to contextualize the request.
4. **Query generation:** The LLM translates the interpreted intent into executable query language, such as DAX or Power Query M.

5. **Validation and execution:** The generated query is syntactically validated and executed within Power BI, applying security constraints such as RLS.

6. **Visualization selection and rendering:** The system determines the most appropriate visualization type and renders the results.

7. **Narrative generation and conversational iteration:** An explanatory narrative is produced, and the system supports follow-up queries, enabling iterative and conversational data exploration.

This pipeline transforms unstructured user input into structured analytical outputs, effectively bridging the gap between natural language and formal query systems.

C. Enterprise Implementation Patterns

The adoption of this architecture can be observed through different implementation patterns, depending on organizational requirements, technical capabilities, and governance constraints:

- **Pattern 1: Native Integration (Power BI Copilot)**
  This approach leverages out-of-the-box capabilities provided by Microsoft, where orchestration, LLM processing, and security are fully integrated. It offers rapid deployment and ease of use but requires licensing through Microsoft Fabric or Power BI Premium/PPU.

- **Pattern 2: Hybrid Integration (Azure OpenAI)**
  In this model, Power BI is primarily used as the visualization layer, while orchestration logic and prompt engineering are developed and managed by the organization using services such as Power Apps, Azure Functions, or custom middleware. This approach provides greater flexibility and customization.

- **Pattern 3: Custom Solution (Open Source or Multi-Cloud)**
  This pattern involves the use of Power BI REST APIs in combination with open-source LLMs or third-party cloud services. It offers maximum control over infrastructure, model selection, and data governance, at the cost of increased implementation complexity and maintenance overhead.

ARCHITECTURE FOR INTEGRATING LLMS INTO BUSINESS INTELLIGENCE

The integration of Large Language Models (LLMs) into Business Intelligence (BI) platforms represents a paradigm shift in how users interact with analytical systems. Rather than relying on predefined queries or manual report construction, users can engage with data through natural language, enabling a more intuitive and accessible analytical workflow. This section describes the core functional components that support this transformation, focusing on the interaction between natural language processing, analytical query generation, and user-centric insight delivery.

One of the foundational capabilities of this architecture is Natural Language to Structured Query Translation (NLP-to-Query). This module interprets user requests expressed in everyday language and transforms them into formal analytical expressions, such as DAX (*Data Analysis Expressions*) or SQL (*Structured Query Language*). The process involves multiple stages, including semantic intent detection, entity recognition, and mapping of linguistic constructs to elements of the semantic data model. For example, temporal expressions such as “last month,” “year-to-date,” or “same period last year” are translated into formal time-intelligence functions supported by the analytical engine. By automating this translation, the system significantly reduces the need for users to understand query syntax, lowering the technical barrier and enabling more inclusive data exploration.

Another critical capability is the synthetic generation of DAX code, where LLMs produce executable analytical expressions based on concise functional descriptions. This functionality is particularly relevant in platforms such as Power BI, where DAX serves as the primary language for defining business logic and metrics. Through prompt-driven generation, LLMs can construct complex measures that incorporate filtering conditions, aggregation logic, and contextual evaluation rules. This accelerates the semantic modeling process, facilitates rapid prototyping of analytical scenarios, and reduces common errors associated with filter and row context misinterpretation. A representative example is the calculation of Month-over-Month (MoM) revenue growth, where the model automatically defines variables for current and prior periods and computes the percentage variation. Such generated expressions adhere to established best practices in BI modeling and can be reused across reports and dashboards, as illustrated in Figure 1.

```
Crecimiento MoM % =
VAR IngresosMesActual = SUM(Ventas[Ingresos])
VAR IngresosMesAnterior =
  CALCULATE(
    SUM(Ventas[Ingresos]),
    DATEADD(Calendario[Fecha], -1, MONTH)
  )
RETURN
  DIVIDE(
    IngresosMesActual - IngresosMesAnterior,
    IngresosMesAnterior,
    0
  )
```

Figure 1. Example of a DAX measure to calculate monthly revenue growth (MoM) in Power BI. Source: Author(s)

Beyond query and code generation, LLM-integrated BI systems provide the capability for generation of explanatory narratives, which complements traditional data visualization. Instead of requiring users to interpret charts and tables independently, the system produces natural language summaries that highlight key findings, identify anomalies, and contextualize results in relation to historical trends or predefined targets. These narratives act as an interpretative layer that bridges the gap between raw data and decision-making, particularly for non-technical stakeholders. Furthermore, the ability to dynamically adapt the tone, level of detail, and terminology of these narratives enhances their usability across different organizational roles, from operational analysts to executive decision-makers.

A key enabler of this interactive paradigm is the maintenance of conversational context, which allows the system to retain memory of previous interactions and support iterative query refinement. Rather than requiring users to restate their full analytical intent in each query, the system leverages contextual continuity to interpret follow-up questions, enabling operations such as drill-down (exploring finer levels of detail) and drill-across (navigating across dimensions). This conversational capability reduces cognitive load, improves analytical efficiency, and fosters a more natural exploration process, closely resembling human dialogue.

Together, these capabilities natural language query translation, automated code generation, narrative synthesis, and conversational context management, form the foundation of LLM-driven augmented analytics. By integrating these components into BI architectures, organizations can move toward more intelligent, adaptive, and user-centric analytical systems that enhance both accessibility and analytical depth.

PRACTICAL IMPLEMENTATION

The practical implementation of Large Language Models (LLMs) within Business Intelligence (BI) environments requires the integration of multiple technological components, including data platforms, orchestration layers, and AI services. Unlike the conceptual architecture, which focuses on functional capabilities, practical implementation addresses how these components are deployed, configured, and integrated within real-world enterprise ecosystems. This section outlines the main implementation approaches, highlighting their architectural characteristics, trade-offs, and operational considerations.

A common starting point for organizations is the adoption of native integration capabilities, particularly through tools such as Power BI Copilot. In this approach, generative AI functionalities are embedded directly within the BI platform, providing out-of-the-box support for natural language querying, automatic report generation, DAX expression synthesis, and narrative creation. This model significantly reduces implementation complexity, as orchestration, model hosting, and security mechanisms are managed within the Microsoft ecosystem. However, it requires specific licensing schemes, such as Microsoft Fabric or Power BI Premium, and may impose limitations in terms of customization and control over model behavior [1].

For organizations requiring greater flexibility, a hybrid implementation approach can be adopted, leveraging services such as Azure OpenAI. In this architecture, Power BI functions primarily as the visualization and data access layer, while an intermediate application layer handles user interaction, prompt engineering, and orchestration logic. Natural language queries are processed by the LLM through API calls, generating analytical instructions such as DAX or Power Query M expressions, that are subsequently executed within the BI environment. Supporting services such as Azure Key Vault and Application Insights are commonly integrated to ensure secure credential management and system monitoring [3], [13]. This approach enables fine-grained control over prompts, model selection, and response validation, making it suitable for organizations with advanced technical capabilities.

A more advanced alternative is the development of custom or API-based architectures, which utilize Power BI REST APIs to enable full programmatic interaction with datasets, reports, and dashboards [12]. In this model, organizations can integrate LLMs hosted on different platforms including open-source or multi-cloud environments, into their analytical workflows. This approach provides maximum flexibility in terms of infrastructure, allowing organizations to tailor the system according to specific requirements related to performance, cost, or data sovereignty. However, it also introduces additional complexity in areas such as orchestration, latency management, and system maintenance.

Within these custom architectures, emerging standards such as the Model Context Protocol (MCP) play a critical role in improving interoperability between LLMs and enterprise systems. MCP defines a structured client–server mechanism through which LLMs can access semantic model metadata, measure definitions, and representative query examples [4]. By enhancing the grounding of model responses, MCP reduces the likelihood of generating syntactically correct but semantically inconsistent outputs, thereby improving reliability and alignment with organizational data models.

Across all implementation approaches, several cross-cutting considerations must be addressed. First, the quality and design of the semantic model particularly in platforms like Power BI, directly influence the accuracy of LLM-generated queries and insights [5]. Second, robust security mechanisms, including Row-Level Security (RLS), authentication, and data governance policies, are essential to ensure compliance with organizational and regulatory requirements [9]. Third, prompt engineering and validation pipelines must be carefully designed to mitigate risks associated with hallucinations or incorrect query generation.

In summary, the practical implementation of LLM-integrated BI systems can be approached through native, hybrid, or fully customized architectures, each offering different levels of control, scalability, and complexity. The selection of an appropriate approach depends on organizational priorities, technical maturity, and governance requirements. When properly implemented, these architectures enable a seamless integration between natural language interaction and analytical processing, unlocking the full potential of augmented analytics in enterprise environments.

FUTURE TRENDS

In summary, the practical implementation of LLM-integrated BI systems can be approached through native, hybrid, or fully customized architectures, each offering different levels of control, scalability, and complexity. The selection of an appropriate approach depends on organizational priorities, technical maturity, and governance requirements. When properly implemented, these architectures enable a seamless integration between natural language interaction and analytical processing, unlocking the full potential of augmented analytics in enterprise environments.

The evolution of Large Language Models (LLMs) integrated into Business Intelligence (BI) platforms is expected to reshape the analytical landscape significantly over the coming years. These advancements are not limited to incremental improvements in user interfaces, but rather point toward a broader transformation in how data is processed, interpreted, and operationalized within organizations.

One of the most prominent trends is the continued progression toward augmented analytics, which integrates machine learning and artificial intelligence techniques to automate and enhance tasks across the entire analytical lifecycle. While current implementations such as the integration of LLMs into Power BI, primarily focus on natural language interaction and automated insight generation, future systems are expected to extend these capabilities toward fully autonomous analytical processes. These include continuous data monitoring, real-time anomaly detection, proactive alerting mechanisms, and conversational predictive analytics. In particular, the incorporation of “what-if” scenario analysis through natural dialogue will allow users to simulate alternative decision paths interactively, reducing the gap between data exploration and strategic planning.

In parallel, a transition is anticipated from traditional dashboard-based interfaces to conversational BI environments, where interaction with data is mediated by intelligent agents. These agents, often specialized by functional domain (e.g., finance, sales, operations), will be capable of executing multi-step analytical workflows, chaining queries, and synthesizing results into executive-level reports. Moreover, they will be deeply integrated into enterprise ecosystems, including Customer Relationship Management (CRM) and Enterprise Resource Planning (ERP) systems, as well as collaboration platforms such as Microsoft Teams and Slack. This shift will redefine the role of BI systems, positioning conversational interfaces as the primary mode of analytical engagement and significantly lowering the barrier to entry for non-technical users.

Another key direction is the emergence of decision intelligence systems, which extend beyond descriptive and predictive analytics to incorporate prescriptive capabilities. Rather than simply presenting data or identifying patterns, these systems will actively recommend actions based on analytical findings. This involves the integration of optimization models, simulation engines, and feedback loops that learn from historical decisions and their outcomes. As a result, organizations will be able to move toward more proactive and adaptive decision-making processes, strengthening a data-driven culture that is accessible across all organizational levels.

Additionally, the development of specialized and fine-tuned language models is expected to play a critical role in improving the effectiveness of LLM-driven BI systems. Instead of relying solely on general-purpose models, organizations will increasingly adopt domain-specific LLMs trained on sector-specific data, such as healthcare, finance, or retail. Furthermore, enterprise-level fine-tuning will allow these models to incorporate internal business rules, semantic definitions, and institutional knowledge. This customization process, combined with continuous learning from user interactions and feedback, will enhance the accuracy, relevance, and contextual alignment of generated insights.

Overall, these trends indicate a shift toward more autonomous, context-aware, and user-centric analytical systems. As LLMs continue to evolve, their integration into BI platforms will not only improve efficiency but also fundamentally redefine how organizations generate, interpret, and act upon data-driven insights.

CONCLUSIONS

The integration of Large Language Models (LLMs) into analytical platforms such as Power BI represents a profound transformation in how organizations access, interpret, and operationalize information. This integration expands the scope of business intelligence and redefines the interaction between users and analytical systems by removing traditional technical barriers.

Key findings include the democratization of data access through natural language interfaces, significant reductions in analysis time, improved communication of results through executive narratives, and architectural flexibility that enables the combination of native and customized solutions.

However, successful adoption requires addressing challenges related to the quality of semantic models, data governance and security, and the realistic management of expectations regarding LLM performance. Robust validation processes and human oversight remain essential, particularly for high-impact decisions.

Organizations that effectively integrate these capabilities can achieve more agile decision cycles, foster innovation in data usage, and optimize analytical talent. The trend points toward conversational and autonomous BI systems, where the distinction between technical and business users becomes increasingly blurred, shifting the focus from syntactic skills to competencies in question formulation and critical thinking.

Ultimately, the integration of LLMs with Power BI is not merely an incremental improvement in user interfaces, but a paradigm shift in the relationship between data and decision-making. Organizations that understand and capitalize on this transformation will be better positioned to thrive in an increasingly data-driven economy.

AUTHOR CONTRIBUTION

The authors' contributions to this article are as follows:
Dagoberto Altamar Pacheco: Conceptualization, research, methodology design, data analysis, visualization, writing, and editing.
Pedro Luis Torres Alvarez: Data analysis, system design, validation, visualization, writing, and editing.

The authors participated in the review of the results and approved the final version of the article.

CONFLICT OF INTERESTS

The authors declare that they have no interests or financial relationships that could have influenced this work.

REFERENCES

[1] Microsoft Corporation, “Power BI Copilot Documentation,” *Microsoft Learn*, 2024. [Online]. Available: https://learn.microsoft.com/en-us/power-bi/create-reports/copilot-introduction

[2] Gartner, Inc., “Magic Quadrant for Analytics and Business Intelligence Platforms,” *Gartner Research*, 2024.

[3] Azure OpenAI Service Team, “Enterprise Integration Patterns for Azure OpenAI,” *Microsoft Azure Documentation*, 2024. [Online]. Available: https://learn.microsoft.com/en-us/azure/ai-services/openai/

[4] Anthropic, “Model Context Protocol Specification,” 2024. [Online]. Available: https://modelcontextprotocol.io/

[5] Microsoft Fabric Team, “Semantic Models in Microsoft Fabric,” *Microsoft Learn*, 2024. [Online]. Available: https://learn.microsoft.com/en-us/fabric/

[6] Power BI Team, “DAX Guide and Best Practices,” *SQLBI*, 2024. [Online]. Available: https://dax.guide/

[7] T. Brown *et al*., “Language models are few-shot learners,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 1877–1901, 2020.

[8] A. Vaswani *et al*., “Attention is all you need,” *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[9] Microsoft Security Team, “Security and Governance in Power BI,” *Microsoft Learn*, 2024. [Online]. Available: https://learn.microsoft.com/en-us/power-bi/enterprise/

[10] Gartner, Inc., “Augmented Analytics: The Future of Data and Analytics,” *Gartner Research Report*, 2023.

[11] A. Ferrari and M. Russo, *The Definitive Guide to DAX: Business Intelligence with Microsoft Excel, SQL Server Analysis Services, and Power BI*, 2nd ed. Microsoft Press, 2023.

[12] Power BI Team, “Power BI REST API Reference,” *Microsoft Documentation*, 2024. [Online]. Available: https://learn.microsoft.com/en-us/rest/api/power-bi/

[13] Azure Architecture Center, “AI-Enriched Business Intelligence Solutions,” *Microsoft Azure Architecture*, 2024. [Online]. Available: https://learn.microsoft.com/en-us/azure/architecture/

[14] R. Kimball and M. Ross, *The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling*, 3rd ed. Wiley, 2013.

[15] OpenAI, “GPT-4 Technical Report,” *OpenAI Research*, 2023. [Online]. Available: https://openai.com/research/gpt-4