



(RESEARCH ARTICLE)



Domain Aware Prompt Engineering

Subhash Taravarthi ^{1,*}, Raghunath Reddy Koilakonda ¹, Venkatasatyaravikiran Bikkavolu ² and Saikrishna Tarakampet ¹

¹ Celina, Texas.

² Louisville, Kentucky.

World Journal of Advanced Engineering Technology and Sciences, 2025, 16(01), 569-574

Publication history: Received on 16 June 2025; revised on 24 July 2025; accepted on 26 July 2025

Article DOI: <https://doi.org/10.30574/wjaets.2025.16.1.1249>

Abstract

Introduction: In this world of Generative AI, it is very important to understand the thought process for generating accurate results. Implementing a Generative AI application using different models has become a very common practice using prompt engineering. Generic Prompting often results in hallucinations, compliance risks and lack of actionable intelligence. To overcome these problems, we introduce Domain Aware Prompt Engineering which is an emerging technique to tailor Gen AI outputs on enterprise specific knowledge, semantics and decision making.

Objectives: In this world we have multiple domains like Finance, Marketing, Supply Chain, Telecom and many more... Providing the insights of each domain in the prompt engineering often gives excellent solutions. Need to understand the cosmetics of each domain and put the flow in your prompt during the implementation of our enterprise Gen AI applications. This article demonstrates how to operationalize domain aware prompt engineering to strategically enable the building of accurate, secure and contextually intelligent Gen AI applications in enterprise environments.

Methods: The proposed architecture caters with multiple functional domains like Finance, Supply Chain, Marketing, Telecom for a Text to SQL implementation where it requires domain specific to understand the functional aspects of underlying data. Give business users an area of self-service implementation on UX/UI where they utilize their domain knowledge to leverage high accuracy. This domain knowledge emphasizes their functional thought process.

Results: The Azure-based Gen AI Text-to-SQL architecture with domain aware prompt engineering has emerged as the enterprise applications. It increased the accuracy up to 95% based on domain knowledge of SMEs, and empowered non-technical users to craft precise SQL queries using natural language, which means less dependence on IT. The ability to seamlessly query across databases like Snowflake, Databricks, and Oracle has really sharpened decision-making. Plus, with Azure AD and RBAC in place, security compliance is a solid 100%. Thanks to Azure services, deployment is scalable and boasts a reliability rate of 99.9%. A case study in the supply chain sector revealed that query times plummeted from days to mere minutes, leading to a 50% boost in analyst productivity and an 80% drop in errors, all of which enhances strategic agility.

Conclusions: A scalable GenAI Text-to-SQL setup that leverages Azure OpenAI, Fast API, and various Azure services is transforming the way enterprises approach Decision Intelligence. It allows users to make natural language queries, gain insights across different databases, and maintain strong governance. This not only lessens the dependency on IT but also enhances agility through effective change management.

Keywords: Generative AI (Gen AI); Large Language Models (LLMS); Domain Knowledge; Domain Expertise; Functional Areas Like Finance; Supply Chain; Billing; Telecom

* Corresponding author: Subhash Taravarthi

1. Introduction

1.1. Problem statement

Enterprises are adopting Generative AI to push business users with natural language for querying structured and unstructured data. One the best use cases is Text-to-SQL for both structured and unstructured data. Usually, users are prompted to ask questions in plain language and generate accurate SQL queries based on the generic prompting techniques.

However, generic prompting techniques often has challenges in enterprise environments, especially different functional domains like supply chain management or financials, where accurate query generation depends on deep understanding of

- Domain-specific terminology (e.g., “lead time,” “overseas supplier,” “last quarter,” “revenue”, “net income”, “gross margin”),
- Enterprise-specific logic and SLAs,
- Structured schema and metadata relationships,
- Unstructured metadata
- Role-based intent and contextual expectations.
- This results in inaccurate, incomplete, misleading results which raises the trust of business often lying back with negative thought processes.

There comes the thought process domain-aware prompt engineering which can

- Align GenAI outputs with enterprise semantics and business logic,
- Profiling the metadata to ensure better semantics of metadata descriptions.
- Reduce hallucination and improve accuracy of generated SQL,
- Increase of accurate results providing better business faith and reliance on the product.

Solving this problem is key to operationalizing trustworthy GenAI systems for real-time decision intelligence in data-driven enterprises.

1.2. Solution approach

To achieve this, we will discuss a few areas of Domain aware prompt engineering with multiple functional areas which can leverage the advantage of prompt engineering to attain accurate results.

Let’s talk about the Supply Chain area where we deal with multiple tables with some calculated metrics like Lead Time and so on.

1.3. Define Domain Scope and Schema

- Identify few tables based on the domain for Supply Chain
- Products dim (productid, name, category)
- Suppliers dim (supplier, name, country, rating)
- Supply_orders_dim (order, productid, supplier, order date, delivery date)
- Time dim (timeid, date, year, quarter, month)
- Inventory fact (productid, location, stock level, reorder point)

Treating this a single table, for better slicing and dicing, create meaningful view *Inventory View* to ensure that we have proper questions answered more functional to achieve business needs. This view is generated from above tables for better slicing and dicing.

1.4. Business Rules

- Make sure any question asked is based out of current year
- Month Format should be YYYYMM only
- Use CTE’s for generating responses
- Some calculated metrics

- Lead Time = delivery date - order date
 - Overseas Supplier = country! = 'USA'
 - Last Quarter = depends on current date; include logic to define it dynamically.
-

2. Define Table / View Metadata for each column involved

As we are aware that we built a view out of all these base tables, to ensure that we are having slice and dice based on multiple columns and provide proper technical descriptions.

- Year: This represents year
- Delivery Date; Date specified based when it is delivered
- Order Date: This is the ordered date
- Country: Represents of each country of the world. For Example; United States as US., India as IN.

2.1. Build a Supply Chain-Aware Prompt Template:

You are a SQL expert in generating SQL queries based on user requests, using the provided schema and business rules and generate with standards and provide proper aliases with standard naming conventions

2.2. Schema

2.2.1. Inventory View

Business Rules

- Lead time = DATEDIFF (day, order date, delivery date)
- Overseas suppliers = country! = 'USA'
- Last quarter = previous 3-month period from today's date

2.3. User Query

2.3.1. "What are the top 5 products?"

SQL Generated

SELECT

- name AS product name,
- COUNT (*) AS top products

2.4. FROM

2.4.1. inventory view

WERE

- year = YEAR(CURRENT_DATE)
- GROUP BY
- name
- ORDER BY

top products DESC

LIMIT 5

2.5. User Query

"What are the top 5 products with longest lead times from overseas suppliers for last quarter"

2.5.1. SQL Generated

WITH lead-times AS (SELECT name AS product_name, DATEDIFF (day, order_date, delivery_date) AS lead-time FROM inventory_view WHERE year = YEAR(CURRENT_DATE))

AND country! = 'USA' AND order_date BETWEEN DATE_TRUNC ('QUARTER', DATEADD (QUARTER, -1, CURRENT_DATE)) AND CURRENT_DATE) SELECT product_name, AVG (lead-time) AS avg_lead_time

FROM

- lead_times
- GROUP BY
- product_name

ORDER BY

- avg_lead_time DESC

LIMIT 5

Let's talk about the Finance area where we deal with multiple tables with some calculated metrics like Lead Time and so on.

2.6. Business Rules

- Active Accounts: status = 'Active'
- High-Value transaction: amount > 10,000
- Domestic Account: country = 'USA'
- Quarter Format: User YYYY-QQ (e.g., 2024-Q2)
- Use CTEs and standard naming conventions in SQL
- Only include current fiscal year (year= YEAR(CURRENT_DATE))

2.7. User Query

"What is the total transaction volume from domestic accounts in the last quarter?"

2.7.1. SQL Generated

SELECT

SUM (amount) AS total_domestic_volume

FROM

financials_view

2.7.2. WERE

country = 'USA'

- AND transaction_date BETWEEN
- DATE_TRUNC ('QUARTER', DATEADD (QUARTER, -1, CURRENT_DATE))
- AND CURRENT_DATE

2.7.3. User Query

"Show top 5 customer segments by number of high-value transactions this year."

2.7.4. SQL Generated

WITH high_value_txns AS (SELECT Segment, transactinide FROM financials_view

WERE

- amount > 10000
- AND status = 'Active'
- AND year = YEAR(CURRENT_DATE)

SELECT

```
segment, COUNT (transactinide) AS high_value_count FROM high_value_txns GROUP BY segment ORDER BY high_value_count DESC LIMIT 5
```

Let's talk about the Telecom industry where we deal with multiple tables with some calculated metrics like Lead Time and so on.

2.8. Business Rules

- Active Accounts: status = 'Active'
- High-Value transaction: amount > 10,000
- Domestic Account: country = 'USA'
- Quarter Format: User YYYY-QQ (e.g., 2024-Q2)
- Use CTEs and standard naming conventions in SQL
- Only include current fiscal year (year= YEAR(CURRENT_DATE))

You are a telecom SQL expert. Use the Telecom View to generate SQL queries based on user input. Incorporate business logic (e.g., heavy data usage, uptime SLAs), use CTEs when needed, and ensure queries reflect current year.

2.9. User Query

"Which regions have the heaviest data users in the last 30 days?"

2.10. SQL Generated

WITH heavy users AS (SELECT region, customer FROM telecom view

2.10.1. WERE

- dating > 50
- AND status = 'Active'
- AND usage date >= DATEADD (DAY, -30, CURRENT_DATE)
- SELECT region, COUNT (DISTINCT customer) AS heavy_user_count
- FROM heavy users GROUP BY region ORDER BY heavy_user_count DESC

2.11. User Query

"Get average network uptime by region for this year and flag regions below SLA"

2.11.1. SQL Generated

- SELECT region, AVG (uptime percent) AS argutite
- FROM telecom view

2.11.2. WHERE

- year = YEAR(CURRENT_DATE) GROUP BY region
- HAVING AVG (uptime percent) < 99

3. Conclusion

Most of the enterprises integrate with business workflows and they say Gen AI doesn't understand business. But it is to bridge a gap on how enterprise data is structured, governed and interpreted. Here domain – aware prompting becomes indispensable. It's not only writing prompts but it's with domain knowledge, operational logic, data semantics and access roles.

The future of AI isn't not the fine tuning the models or use expensive deployments, rather it is all about how we teach models to speak our language with domain awareness.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

- [1] OpenAI. (2023). OpenAI Cookbook: Prompt Engineering Guide. Retrieved from <https://github.com/openai/openai-cookbook>
- [2] LangChain. (2024). LangChain Documentation. Retrieved from <https://docs.langchain.com>
- [3] Yu, T., et al. (2018). Spider: A Large-Scale Human-Labeled Text-to-SQL Dataset for Complex and Cross-Domain Queries. Retrieved from <https://yale-lily.github.io/spider>
- [5] Zhong, V., Xiong, C., & Socher, R. (2017). Seq2SQL: Generating Structured Queries from Natural Language using Reinforcement Learning. arXiv:1709.00103
- [6] Snowflake. (2023). Understanding Access Control and Role Hierarchies. Retrieved from <https://docs.snowflake.com>
- [7] Databricks. (2024). Architecting GenAI Systems with Enterprise Data Lakes. Retrieved from <https://www.databricks.com/blog>