CHAPTER 1
INTRODUCTION

1.1 Background

Kitabisa.com, as the largest donation platform in Indonesia, manages a very large volume of data every day, ranging from donation data and social-campaign activity to payment transactions. All of this data is consolidated into an enterprise data warehouse to ensure that business units use a consistent source of data (Ponnusamy, 2023). Although centralized data storage infrastructure is already available, fast access to information remains a major challenge for non-technical units, especially operational teams whose data needs are highly dynamic and directly support daily operations.

In operational work, teams face a pace that requires fast and adaptive decision-making. Information needs often appear repeatedly and in an ad hoc manner, such as monitoring donation trends or detecting anomalies in campaigns in real time. However, independent data exploration is limited because structured data architecture requires SQL query skills and an understanding of complex database schemas (Quamar et al., 2022; Sura, 2025). This technical barrier forces operational teams to rely on manual data-request flows or ticketing systems, which can slow down decision-making.

This research designs and implements a retrieval-augmented generation (RAG)-based conversational AI agent that enables operational teams to ask data-related questions in natural language and receive answers directly from the enterprise data warehouse. This approach was selected to reduce schema-mapping errors and keep each answer aligned with the company's business documentation (Khan et al., 2025).

This research focuses on the AI agent's ability to run a controlled workflow, starting from question clarification, SQL construction, and result validation before the answer is presented to users. By using the enterprise data warehouse as the

single source of truth, the system keeps each answer aligned with standardized business metric definitions at Kitabisa.com. The same single-source approach also supports stricter security controls, including read-only access and sensitive-data masking for each interaction.

The data in this research is Business Intelligence (BI) data that is already available in the Kitabisa.com environment. The data comes from transaction logs and operational databases that are integrated and consolidated into the company data warehouse. The scope includes contribution data or Gross Written Premium (GWP), claim data, policy data, participant data, insurance product data, payment methods, membership status, and other operational data used for reporting and business analysis. In practice, business users often need this information through operational questions such as "last month's GWP", "refund claims", "SJK policies", "tabarru", or "BSI production".

However, the business terms used by users do not always have direct equivalents that are easy to identify in data warehouse tables and columns. As a result, data exploration still often requires help from the BI or data team to translate user needs into the correct SQL query. This research therefore designs and implements a Retrieval-Augmented Generation (RAG)-based AI agent as an additional interaction layer on top of existing BI data. The implemented system is not intended to replace BI work. Instead, it helps non-technical users access data through natural language, maps business terms to BI metadata, generates validated SQL queries, and presents answers in a faster, safer, and more traceable way.

To maintain security during initial development and evaluation, this research uses a synthetic data warehouse that represents operational patterns for donations, campaigns, policies, claims, insurance products, payment methods, and settlement. Synthetic data is used to test the Text-to-SQL flow, query validation, sensitive-data masking without exposing confidential production data.

1.2 Problem Formulation

Based on this background, the research questions are formulated as follows:

1) How can a RAG-based conversational AI agent be designed to help non-technical users access information from an enterprise data warehouse through natural language?
2) How can business-term mapping, metadata retrieval, SQL generation, query validation, and answer presentation be implemented so that the results remain accurate and traceable?
3) How can SQL security validation, read-only access, statement timeout, and sensitive-data masking be applied to reduce the risk of data leakage?
4) How can the AI agent be evaluated in terms of Text-to-SQL accuracy, execution success, dialog quality, response latency, task-completion efficiency, and Usability Testing?

1.3 Research Scope

To keep the research scope focused, the limitations of this study are defined as follows:

1) The main data source is limited to an enterprise data warehouse or a synthetic data warehouse that represents the company analytical structure as a single source of truth for query execution.
2) The system supports only read-only analytical queries and does not cover write operations such as INSERT, UPDATE, DELETE, ALTER, DROP, or CREATE.
3) Data access is directed through curated analytical views and semantic metadata. Detailed network or infrastructure security discussion is outside the scope of this research.
4) This research focuses on the design and prototype of a RAG-based AI agent, SQL validation, result masking and limited evaluation on business-question scenarios.
5) Fine-tuning large language models and developing model architecture from scratch are not discussed in depth.

1.4 Objectives and Benefits

1.4.1 Objectives

The objectives of this research are as follows:

1) Produce the architecture design and prototype of a retrieval-augmented generation conversational AI agent that is integrated with a data warehouse to provide self-service data access through natural-language conversation.
2) Implement an AI agent workflow that includes authentication, semantic metadata retrieval, business-term mapping, Text-to-SQL, and answer presentation so that each result can be traced.
3) Apply SQL security validation, read-only access, statement timeout, and sensitive-data masking to reduce the risk of data leakage.
4) Evaluate the AI agent performance in processing operational-question scenarios by measuring query accuracy, execution success, response latency, task-completion efficiency, and Usability Testing.

1.4.2 Benefits

The expected benefits of this research are as follows:

1) For the Kitabisa.com operational team, this system provides faster and more independent access to information without requiring users to write SQL or understand the detailed data warehouse schema.
2) For the Kitabisa.com data team, this system helps reduce repetitive ad hoc data requests so that work capacity can be redirected to data quality improvement, metadata enrichment, and data governance.
3) For academic institutions, this research provides a case study on the implementation of RAG, Text-to-SQL, security validation, and synthetic data warehouses in an enterprise environment.