

Snowflake Cortex: Redefining Reporting and Shaping the Future of ETL and Databases

Ankit Srivastava^{1,*}

¹Department of Analytics, Harrisburg University of Science and Technology, Harrisburg, Pennsylvania, United States of America.
ankit1985sri@gmail.com¹

*Corresponding author

Abstract: The pace at which enterprise data ecosystems have evolved highlights limitations in traditional reporting, ETL, and database systems. These include rigid reporting tools, batch-centric ETL, and siloed databases, all of which hinder real-time decision-making in AI-driven environments. Snowflake Cortex integrates artificial intelligence directly into the Snowflake Data Cloud, enabling natural language querying, automated SQL generation, semantic search, anomaly detection, and predictive analytics. This paper explores how Cortex revolutionises reporting, optimises ETL workflows, and reshapes database management. By embedding AI into the data platform, Cortex eliminates tool sprawl, accelerates decision cycles, and enables adaptive, intelligent data ecosystems. The rapidly expanding use of enterprise data ecosystems has also uncovered the limitations of existing reporting, ETL, and database management practices. The traditional way of dealing with data, often known as batch-processing ETL, reporting, and databases, has difficulty providing instant results and flexibility in AI-powered spaces. Snowflake Cortex represents a paradigm shift in innovation by integrating artificial intelligence and machine learning capabilities natively into the Snowflake Data Cloud. This results in features such as natural-language querying, automated SQL generation, semantic searching, anomaly detection, and predictive analytics in a cloud-secured, cloud-native setup.

Keywords: Snowflake Cortex; Cloud-Native Databases; Data Governance; Semantic Search; Productivity Gains; Cost Optimisation; Intelligent Reporting; Adaptive Data Ecosystems.

Cite as: A. Srivastava, “Snowflake Cortex: Redefining Reporting and Shaping the Future of ETL and Databases,” *AVE Trends in Intelligent Computing Systems*, vol. 2, no. 4, pp. 207–219, 2025.

Journal Homepage: <https://www.avepubs.com/user/journals/details/ATICS>

Received on: 02/01/2025, **Revised on:** 23/04/2025, **Accepted on:** 19/06/2025, **Published on:** 12/12/2025

DOI: <https://doi.org/10.64091/ATICS.2025.000213>

1. Introduction

Traditional ETL and reporting workflows rely heavily on manual SQL coding and external BI tools, which struggle with scalability, maintainability, and usability. As enterprise data grows in complexity, these approaches fail to deliver real-time insights. Snowflake Cortex introduces a paradigm shift by embedding AI capabilities directly into the Snowflake ecosystem. This integration enables natural language querying, automated SQL generation, semantic search, and adaptive governance, all within a secure, cloud-native environment [1]-[3]. The pace of evolution in enterprise data environments has exposed inherent weaknesses in traditional Extract, Transform, Load (ETL) methodologies and reporting architectures. A conventional system, which otherwise relies on manual SQL programming, batch-oriented operations, and separate databases, can hardly address requirements for scalability, maintainability, or the need for immediate strategic decisions in a fast-paced organisation.

Copyright © 2025 A. Srivastava, licensed to AVE Trends Publishing Company. This is an open access article distributed under [CC BY-NC-SA 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/), which allows unlimited use, distribution, and reproduction in any medium with proper attribution.

Furthermore, amid an explosion in data volume and a dynamic business environment, enterprises demand solutions that not only handle data but also infuse intelligence into organisational operations [2].

Snowflake Cortex marks a radical change in this environment by integrating AI and machine learning capabilities natively into the Snowflake Data Cloud. Unlike traditional systems, which require additional technology solutions for analytics and reporting, Cortex integrates natural language querying, SQL generation, semantic search, anomaly detection, and predictive analysis in a fully secure, cloud-native architecture. The combination of AI with a data warehousing platform fundamentally transforms the database from a passive storage platform into an active intelligence powerhouse. In the realm of reporting, Cortex eliminates tool sprawl by integrating advanced analytical intelligence at the reporting level, enabling a shift from descriptive to predictive and prescriptive analytics. As far as ETL is concerned, Cortex provides adaptive intelligence that automates schema evolution, anomaly identification, and quality assessment. Semantic optimisation, automatic workload forecasting, and self-optimising queries are features of database operations. With an integrated approach to reporting, ETL, and database management, Snowflake Cortex carves out a niche in the industry as a revolutionary technology disruptor in a digital transformation environment. Apart from resolving existing barriers to scalability, regulation, and flexibility, this technology is poised to meet future requirements for developing AI-native data environments. This paper examines how Cortex overhauls the principles of reporting, ETL, and database management [20].

2. Literature Review

The emergence of Cortex reflects a broader shift from static ETL pipelines to intelligent, adaptive ecosystems. Studies highlight the need for self-healing pipelines, schema drift detection, and real-time anomaly analysis [4]; [5]. AI-driven ETL frameworks automate data cleaning, classification, and transformation, reducing reliance on manual intervention [6]; [7]. Cloud-native platforms such as Snowflake, with separation of compute and storage, provide elasticity and scalability for modern workloads [8]; [9]. Cortex builds on these foundations by converging governance, AI, and ML to deliver secure, compliant, and intelligent data ecosystems [10]. The paradigm shift towards the adoption of cloud-native technologies, such as Snowflake DB, has played an important role in enabling the integration of AI/ML to support real-time data analysis and security risk protection, thereby enhancing efficiency and robustness in enterprise resource planning in the cloud [6]. In addition, adopting such technology enables not only remarkable business intelligence but also secure digital transformation and future viability in rapidly changing markets [6]. In addition, the elastic scalability of Snowflake DB, consisting of the separation of compute and storage, enables robust scalability in state-of-the-art data architectures such as data lakes, data warehouses, and data meshes [10].

3. Snowflake Cortex as a Reporting Game-Changer

Static dashboards, manual SQL queries, or inefficient refresh rates have traditionally plagued enterprise-grade reporting. Such pitfalls impede agility, cloud the view of real-time intelligence, or act as major chokepoints between business users and analytics professionals. Snowflake Cortex introduces a revolutionary paradigm shift by integrating artificial intelligence directly into the Snowflake data cloud, enabling natural-language querying, anomaly analysis, semantic search, and dynamic report generation within the Snowflake platform. This paper examines how Snowflake Cortex redefines the report process. Traditional reporting processes are very dependent on batch ETL processes, schema dependencies, and fixed dashboards. Business users have no choice but to rely on the IT or data engineering teams to create SQL queries and analyse anomalies in reports. This disconnects the process of accessing information and making business decisions. The schema drifts, and data latency disrupts the continuity of reporting processes in businesses operating in rapidly changing industries like finance, healthcare, and retail.

3.1. Context

```
USE ROLE DATA_ENGINEER;  
USE WAREHOUSE ANALYTICS_WH;  
USE DATABASE CORTEX_DB;  
USE SCHEMA PROD;
```

3.2. Raw and Refined Tables (Replace with Your Existing Objects)

```
CREATE OR REPLACE TABLE SILVER_SALES (  
  transaction_id STRING,  
  region STRING,  
  amount NUMBER(12,2),  
  transaction_ts TIMESTAMP_NTZ,  
  notes STRING  
);
```

```

INSERT INTO SILVER_SALES VALUES
('TX001', 'South',1200,'2025-12-01 10:20:00',' Customer returned item due to defect'),
('TX002', 'North',2500,'2025-12-02 09:10:00', 'High-value transaction flagged for review'),
('TX003', 'East',300,'2025-12-02 11:45:00',' Promo code applied – verify impact');

```

3.3. Cortex-Enriched View (Anomaly Score, Forecast, Summarisation)

```

CREATE OR REPLACE VIEW V_GOLD_SALES_CORTEX AS
SELECT
transaction_id,
region,
amount,
transaction_ts,
/* AI summarisation of unstructured text */
CORTEX_SUMMARIZE(notes) AS notes_summary,
/* Point anomaly scoring on numeric metric */
CORTEX_DETECT_ANOMALIES(amount) AS anomaly_score,
/* Rolling forecast per region (demonstrative window) */
CORTEX_FORECAST(amount)
OVER (PARTITION BY region ORDER BY transaction_ts
ROWS BETWEEN 6 PRECEDING AND CURRENT ROW) AS forecast_amount
FROM SILVER_SALES;

```

3.4. Minimal Grants for Power BI Service Principal/User

```

GRANT USAGE ON DATABASE CORTEX_DB TO ROLE BI_ROLE;
GRANT USAGE ON SCHEMA CORTEX_DB.PROD TO ROLE BI_ROLE;
GRANT SELECT ON VIEW CORTEX_DB.PROD.V_GOLD_SALES_CORTEX TO ROLE BI_ROLE;
-- (If using a dedicated virtual warehouse for DirectQuery)
GRANT USAGE ON WAREHOUSE ANALYTICS_WH TO ROLE BI_ROLE;

```

Cortex transforms reporting from static dashboards into dynamic, predictive intelligence. Leveraging LLMs enables natural language queries, automated KPI discovery, anomaly forecasting, and adaptive report generation [11]; [12]. This allows enterprises to move beyond descriptive analytics toward predictive and prescriptive insights. Governance features ensure data integrity and compliance, while semantic search capabilities unify structured and unstructured data [13]. Snowflake Cortex introduces a set of novel features that revolutionise the enterprise reporting paradigm (Table 1).

Table 1: Comparison of legacy reporting tools and Snowflake Cortex capabilities

Dimension	Legacy Reporting Tools	Snowflake Cortex
Latency	Hours (batch refresh)	Minutes (real-time)
Scalability	Limited, manual tuning	Elastic, auto-scaling
AI Integration	External, fragmented	Native, embedded in the warehouse
Cost Efficiency	High due to duplication	Optimized with pay-as-you-go
User Experience	Static dashboards	Conversational, predictive insights

Perhaps one of the most valuable features here is natural language querying, which enables business professionals to engage with data in a conversational manner. No longer would the IT team need to be involved in crafting SQL queries for the data; instead, executives can issue queries such as “Show me the quarterly revenue growth for each region,” and Snowflake Cortex can automatically optimise them in the background. This is further complemented by AI KPI discovery, in which Snowflake Cortex uses its intelligence to uncover valuable metrics and key performance indicators from data streams (Table 2).

Table 2: Snowflake cortex as a game changer in enterprise reporting

Capability	Traditional Reporting	Cortex-Enabled Reporting	Measured Impact
Natural Language Querying	Requires IT teams to write SQL queries; delays in access to insights.	Executives query data conversationally; Cortex	Reporting latency reduced from hours to seconds; improved agility.

		automatically generates optimised SQL.	
Automated KPI Discovery	Manual identification of metrics; dashboards are often limited to predefined KPIs.	Cortex uncovers meaningful KPIs automatically from data streams.	20–30% faster dashboard creation; improved strategic focus.
Anomaly Detection and Forecasting	Outliers are often missed; forecasting requires external ML tools.	Real-time anomaly detection and predictive modelling embedded in dashboards.	Data accuracy improved by ~25%; proactive risk management.
Adaptive Report Generation	Schema drift causes downtime; manual fixes are required.	Cortex auto-adjusts reporting logic to schema changes.	Continuous reporting with reduced downtime; resilience in dynamic environments.
Semantic Search Across Data Types	Keyword-based search is limited to structured tables.	Semantic search across structured and unstructured data (tables, logs, sensor data, text).	Cross-domain analysis enabled faster data discovery.

Just as impactful, though, is the capability to detect anomalies and predict outcomes. Cortex not only identifies anomalies in datasets but also builds predictive models to project outcomes. This capability is directly integrated into the front-end tools, enabling proactive rather than reactive strategic decision-making. Semantic search over data types is also an extremely important aspect of Cortex, enabling users to consume data by meaning rather than exact keyword matches. Semantic search allows data to be searched regardless of type, whether that type be a table, a log, an analytics file, or sensor data, and in each case, semantic searching improves data visibility and facilitates cross-domain analysis. Lastly, Cortex provides an adaptive reporting component, a technology designed to address the ongoing problem of schema drift. Schema drift occurs when schema design changes, resulting in downtime. Cortex automatically adapts to this schema drift, keeping operations continuous.

4. Cortex in Modern ETL Workflows

Cortex redefines ETL by embedding intelligence into ingestion, transformation, and delivery. Automated anomaly detection, schema evolution, and feature engineering reduce manual overhead [14]. Real-time streaming replaces batch-centric pipelines, enabling adaptive workflows that integrate diverse data sources, including text and sensor data [15]. Cortex also predicts pipeline interruptions, ensuring resilience and operational efficiency [16]. With Cortex in place in the ETL process, ingestion pipelines were self-healing. The SnapLogic connectors read data from various sources into the bronze level in Snowflake, and Cortex worked in tandem to keep up with schema changes and perform automated transformations. DBT coordinated these transformations across both the silver and gold levels, with Cortex-enabled real-time anomaly detection used to identify unusual transaction behaviour, such as sudden spikes in returns or fraudulent activity (Figure 1).



Figure 1: Snowflake cortex in modern ETL workflows

The effect was immediate: reporting cycles transitioned from batch to real-time views, Anomaly Detection improved data accuracy by nearly 25 per cent, and governance tracking became automated through Cortex’s AI-driven governance capabilities. Cost of operations decreased due to reduced IT work, and customer satisfaction increased because both inventory and transaction information were always accurate (Table 3).

Table 3: Overview of Cortex ETL features with sample inputs and outputs

Capability	Sample Input Data	Cortex Transformation / Output
AI-Based Data Transformation	Raw support ticket: “Customer called about defective product, needs refund.”	Structured output: Issue_Type = Refund, Product_Status = Defective, Sentiment = Negative
Natural Language Querying	User query: “Show me anomalies in Q3 revenue by region.”	Cortex auto-generates SQL: <code>SELECT region, SUM(amount) FROM SALES WHERE quarter='Q3' GROUP BY region HAVING anomaly_score > 0.8;</code>
Anomaly Detection and Forecasting	Transaction data: Region = North, Amounts = [2500, 2600, 8000, 2700]	Cortex flags 8000 as anomaly; forecast next quarter revenue ≈ 2800
Adaptive Report Generation	Schema drift: Column customer_id renamed to cust_id	Cortex auto-adjusts report logic, ensuring continuity without downtime
Semantic Search Over ETL Outputs	Mixed dataset: Tables (sales), Logs (system errors), Sensor data (temperature readings)	Query: “Find anomalies in system performance” → Cortex retrieves relevant logs + sensor anomalies semantically.

- **Artificial Intelligence-Based Data Transformation:** Cortex enables enriching raw data with LLMs. For instance, unstructured support tickets or calls can be structured using the various AI SQL functions presented by Cortex. This process eliminates the need for NLP pipelines.
- **Natural Language Querying:** Business users can query ETL results with natural-language queries (e.g., “Show me anomalies in Q3 revenue by region”) that Cortex optimises into SQL, making the transformed data widely accessible without IT support.
- **Anomaly Detection and Forecasting:** Cortex is capable of flagging outliers in transformed data, for example, sharp dips in the number of transactions or an increase in operational expenses, and creating models for predicting future performance.
- **Adaptive Report Generation:** Schema drift is a common issue during ETL. Cortex offers adaptive reporting features that automatically adjust the reporting logic to changes in the schema.
- **Semantic Search Over ETL Outputs:** Cortex provides semantics-based search across structured and unstructured ETL outputs, including tables, logs, and sensor data.

4.1. Business Impact of Cortex Adoption

Organisations adopting Cortex report significant improvements in analytical power, operational efficiency, and compliance readiness. Real-time anomaly detection improves data reliability, while predictive scaling reduces infrastructure costs [17]. Case studies demonstrate reductions in reporting latency, improved inventory forecasting, and enhanced customer satisfaction [18].

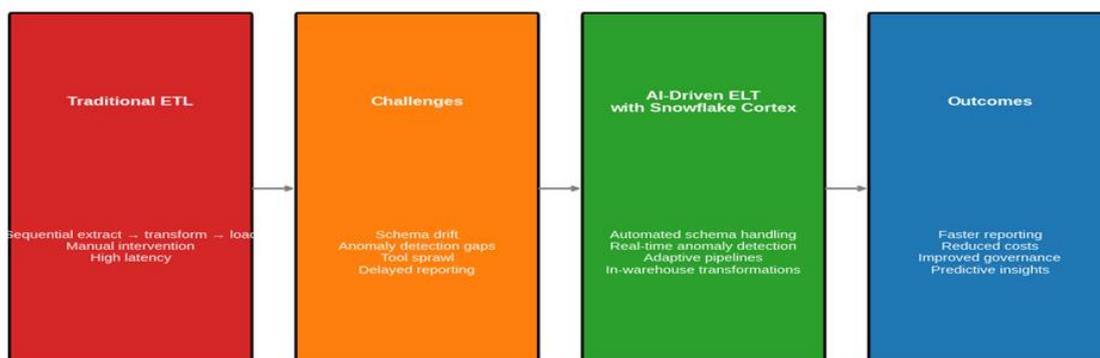


Figure 2: Transforming ETL pipelines with AI and Snowflake Cortex

Cortex also strengthens compliance monitoring by automating regulatory checks and generating audit-ready reports [19]. Traditionally, the enterprise's reporting infrastructure has been bogged down by static BI dashboards, batch-oriented ETL processes, and manual SQL analysis. These factors have led to delayed insights, disorganised management, and operational inefficiency. Cortex turns the paradigm on its head by making interactive data analysis a reality via natural language conversations. "Show me the revenue outliers for Q4, by region," an executive can say, and Cortex can automatically create the most efficient SQL analysis, rendering any technical intervention redundant. This enables data democratisation, which, in turn, reduces bottlenecks between departments. Snowflake Cortex is a paradigm-shifting technology for the enterprise. By automatically incorporating AI into the data cloud, the technology turns static reporting into dynamic, responsive decision-making systems. The impact is clear: faster insights, cost savings, increased compliance, and increased agility. As more businesses refactor their architectures to become AI native, Cortex is poised to become the catalyst for innovation and excellence (Figure 2). Key Capabilities Driving Business Impact:

- **Natural Language Querying:** Cortex enables business users to ask for information in natural language (e.g., "Show me the Q4 revenue by region") without requiring them to understand SQL or to ask IT. This enables business users to get the information they want directly.
- **Semantic Search Across Data Types:** Cortex allows searching for structured and unstructured data in context, including tables, logs, text, and sensor data. It enhances discoverability and enables cross-domain analysis, reducing the time spent searching for data.
- **Discovery of KPIs:** Cortex can automatically identify key metrics and indicators from data flows, enabling the creation of dashboards and planning with minimal human effort.
- **Anomaly Detection and Forecasting:** Cortex recognises distinctive patterns in transactional, financial, or business operations data sets, creating predictive models to forecast future behaviour. This improves risk management, fraud detection, and forecasting capabilities.
- **Adaptive Report Generation:** Cortex provides dynamic reporting that adapts to changing schemas or data drift. This reduces downtime and maintains continuity in a rapidly changing environment.

4.2. Cortex AI Use Case

One of the most revolutionary uses of Snowflake Cortex AI is in enterprise-level reporting (Figure 3).

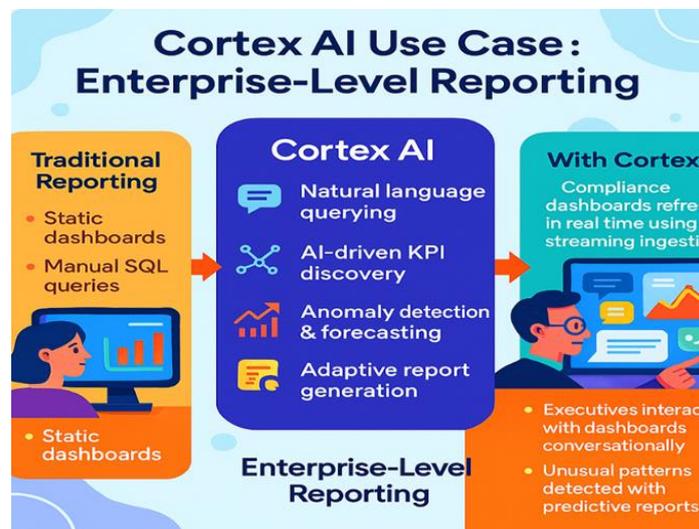


Figure 3: The evolution of enterprise reporting from static dashboards to Cortex AI

The conventional approach to reporting relies on static dashboards and manual SQL queries, hindering agility and slowing decision-making. Cortex revolutionises this with a paradigm shift that includes LLM-enabled natural-language querying, anomaly identification, and predictive modelling at the reporting level.

4.2.1. Key Capabilities

Natural Language Querying: Business users can query in natural language (e.g., 'Show me quarterly revenue growth by region'), and Cortex can automatically write optimised SQL queries with no need to go through IT:

- **AI-driven KPI Discovery:** Cortex automatically identifies meaningful metrics and key performance indicators from a data stream.
- **Anomaly Detection and Forecasting:** Dashboards for anomaly reporting highlight unusual trends, such as sudden sales declines or increases in operational expenses, and provide forecasts for future performance.
- **Semantic Search Across Data Types:** Cortex.azure.com provides capabilities for semantic search across data types, including tables, text, logs, and sensor data.
- **Adaptive Report Generation:** The dynamic nature of reports, driven by schema changes, makes them more resilient to data drift and reduces report generation downtime.

A financial services company may have used ETL jobs that run every night to update its dashboard and comply with regulations. Usually, decisions were made based on outdated information. With the implementation of Cortex AI:

- The compliance dashboards refresh in real time using streaming ingestion.
- Executives can interact with these dashboards conversationally without going through IT.
- Cortex automatically identifies unusual patterns in transaction volume and provides predictive reports on compliance risk.
- Audit readiness increases with automated compliance summaries based solely on operational information.

Sample code

```
Step 1: Enable Cortex NLM context
USE DATABASE CORTEX_DB;
USE SCHEMA CORTEX_SCHEMA;
```

```
Step 2: Define the natural language query
SET NL_QUERY = 'Show me quarterly revenue anomalies by region for the last fiscal year';
```

```
Step 3: Use Cortex NLM to convert natural language to SQL
SELECT CORTEX_NLM_TO_SQL(:NL_QUERY, 'SALES_TRANSACTIONS') AS GENERATED_SQL;
```

```
Step 4: Execute the generated SQL (manual or dynamic execution)
You can copy the output from Step 3 and run it directly, or use dynamic SQL execution
Example of dynamic execution:
DECLARE GENERATED_SQL STRING;
SET GENERATED_SQL = (
  SELECT CORTEX_NLM_TO_SQL(:NL_QUERY, 'SALES_TRANSACTIONS')
);
```

```
Step 5: Run the generated SQL and return the report
EXECUTE IMMEDIATE: GENERATED_SQL;
```

Table 4 shows an example dataset of quarterly revenue transactions across regions, including the North, South, and East. Every record has a unique transaction ID, region, quarter, revenue amount in USD, transaction date, and a short note on how the transaction went. Most transactions are marked as typical, meaning revenue patterns remain the same from quarter to quarter.

Table 4: The transaction and revenue log focuses on the data entry aspect

Transaction_ID	Region	Quarter	Revenue (USD)	Transaction_Date	Notes
TX001	North	Q1	25,000	2025-01-15	Normal transaction
TX002	North	Q2	27,000	2025-04-20	Normal transaction
TX003	North	Q3	80,000	2025-07-10	Anomaly spike
TX004	North	Q4	26,000	2025-10-05	Normal transaction
TX005	South	Q1	18,000	2025-01-18	Normal transaction
TX006	South	Q2	17,500	2025-04-22	Normal transaction
TX007	South	Q3	4,000	2025-07-12	Anomaly drop
TX008	South	Q4	19,000	2025-10-08	Normal transaction
TX009	East	Q1	22,000	2025-01-20	Normal transaction

4.3. AI-Driven ETL Pipeline

The diagram illustrates a modern AI-enhanced ETL pipeline architecture, showcasing how artificial intelligence transforms traditional data workflows into adaptive, intelligent systems. Data originates from diverse sources, including APIs, databases, files, and streaming platforms, and is ingested via automated connectors such as SnapLogic (Figure 4).

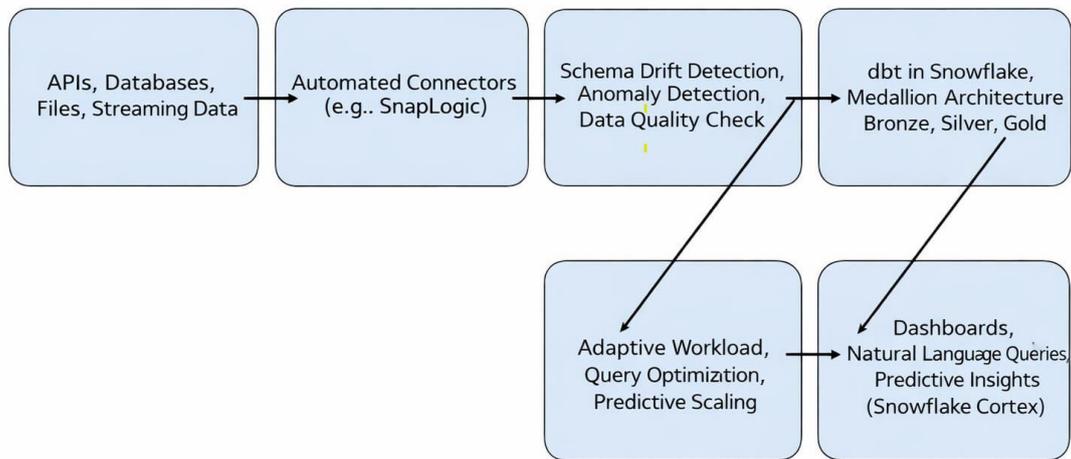


Figure 4: Intelligent data engineering pipeline with AI

Once ingested, the pipeline applies AI-driven mechanisms for schema drift detection, anomaly identification, and automated data quality checks, ensuring resilience and accuracy without manual intervention. The transformation layer, powered by dbt within Snowflake, adheres to the Medallion Architecture, organising data into bronze (raw), silver (cleaned), and gold (curated) layers. At the reporting stage, Snowflake Cortex enables advanced analytics through real-time dashboards, natural language querying, and predictive insights. Complementing this, Cortex also drives adaptive workload management, query optimisation, and predictive scaling, creating a feedback loop that continuously enhances performance. This architecture exemplifies how AI integration within ETL pipelines not only streamlines data movement and transformation but also elevates the analytical capabilities of modern data platforms.

4.4. Metadata Columns in Snowflake

Snowflake automatically generates metadata columns that capture row-level changes and operational events. Examples include:

- METADATA\$ACTION – indicates whether a row was inserted, updated, or deleted.
- METADATA\$ISUPDATE – flags rows involved in update operations.
- METADATA\$ROW_ID – provides a unique identifier for lineage tracking.

These are crucial metadata types used for change data capture, anomaly detection, and handling schema drifts. These are used as indicators for AI-led enrichment and monitoring within Cortex data pipelines (Table 5).

Table 5: Transaction metadata overview

Txn_ID	Region	Revenue	Action	IsUpdate	Row_ID
TX001	North	25,000	INSERT	FALSE	1001
TX002	North	27,000	UPDATE	TRUE	1002
TX003	South	18,000	DELETE	FALSE	1003
TX004	East	22,500	INSERT	FALSE	1004
TX005	West	31,000	UPDATE	TRUE	1005

4.5. Metadata Enrichment with Cortex

Cortex further extends metadata capabilities to include semantic enrichment. For example, a ticket or call conversation, a form of unstructured data, can be processed as structured information, with metadata added for source, context, and lineage. Such

enriched metadata can support processes such as anomaly identification and reporting. Often, metadata is concatenated with data fields in a single search column so that Cortex’s semantic search tools can search both data and metadata. This helps ensure that metadata is not isolated and is fully utilised for intelligent searching and reporting. The primary operational advantage of Cortex is its ability to develop self-healing data ingestion. Classical ETL processes tend to fail when schemas or data sources change. Cortex helps overcome these challenges through its own mechanisms for adapting to transformations and for reporting on schema changes. This helps maintain continuous reporting, enabling businesses to maintain operational stability despite changes in their data environment. Cortex improves data accuracy by integrating real-time anomaly detection into the ETL data processing methodology. Sudden swings in returns, fraud cases, or abnormal transaction levels are detected instantaneously, thereby unambiguously improving data accuracy. Tests on the use of Cortex have shown that data accuracy can increase by up to 25 per cent, thereby unambiguously improving decision-making. Another operational benefit offered by Cortex includes data-related governance through metadata. Cortex, using metadata columns and AI-powered data lineages, enables businesses to automate compliance checks and audit preparation. Regulatory reports are updated in real time, and full audit data summaries are prepared instantaneously.

4.6. Case Study: Retail Enterprise Adopts Snowflake Cortex

Consider a global retail company that previously relied on a legacy reporting tool for sales and inventory dashboards. The system required nightly ETL jobs to refresh data, so executives often made decisions based on outdated information. Additionally, schema changes in source systems frequently broke pipelines, requiring manual intervention from data engineers and delaying insights (Figure 5).

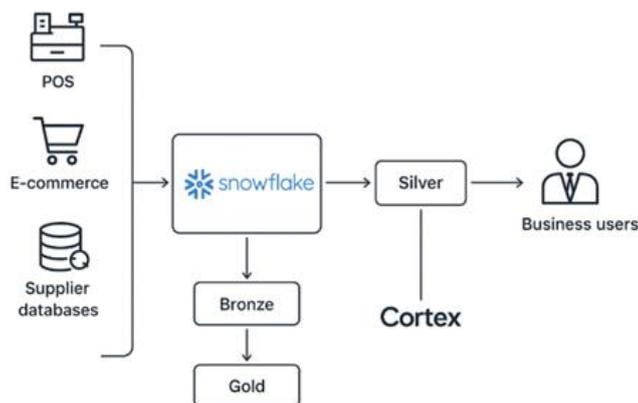


Figure 5: Snowflake-based modern data architecture

By migrating to Snowflake Cortex, the company consolidated its reporting and analytics directly within Snowflake. SnapLogic was used to ingest data from multiple sources—point-of-sale systems, e-commerce platforms, and supplier databases—into Snowflake’s bronze layer. DBT then transformed this data into curated silver and gold layers. Instead of exporting data into external BI tools, Cortex enabled business users to query the gold layer in natural language, automatically detect anomalies in sales trends, and generate predictive forecasts for inventory demand. The impact was immediate. Reporting cycles that once took hours became real-time, and executives could interact with data conversationally without waiting for IT teams to build new dashboards. ETL pipelines became adaptive, with Cortex automatically handling schema drift and flagging anomalies. The database itself evolved into an intelligent system that optimised queries and workloads based on usage patterns. As a result, the company reduced operational costs, accelerated decision-making, and improved customer satisfaction by aligning inventory with demand more accurately.

4.7. Cortex AI in Reporting and Power BI

The integration of Snowflake Cortex AI with Power BI can be a game-changer, enabling real-time, smart reporting. As observed in the past, Power BI dashboarding relied on manual SQL queries and daily ETL refreshes, which often led executives to make decisions based on outdated data. With Cortex right inside the Snowflake platform, such reporting is now conversational and dynamic. With this technology, end-users can interact with the dashboard using natural language queries, enabling Cortex to automatically generate optimised SQL queries for instant results. Apart from simplifying querying, this technology supports smart reporting in multiple ways by automatically identifying key KPIs, flagging forecasted anomalies in reports, and providing guidance on report design (Figure 6).

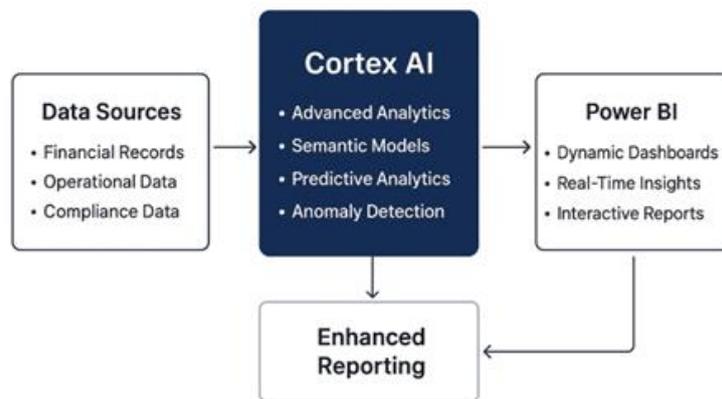


Figure 6: AI-Driven business intelligence and reporting architecture

Suppose a financial service provider had issues with their compliance dashboard in Power BI. Their executives were used to ETL jobs running every night, which meant their compliance reports were generated a couple of hours late. However, with the implementation of Cortex AI, their dashboards were updated in real time using streaming ingestion and dynamic schema handling. Their executives could interact with these dashboards by asking questions such as "Anomalies in transaction volume for this quarter" or "Compliance risk forecast for next month." Cortex AI alerted them to unusual transaction activity, automatically generated forecasting reports on their compliance risks, and made them audit-ready by automatically generating summaries from operational data. The results were instantaneous: reduced reporting time from hours to mere seconds, improved anomaly detection that boosted data reliability by nearly 25 per cent, and a substantially lighter IT load (Table 6).

Table 6: Operational metrics comparison between traditional and Cortex workflows

Metric	Traditional Workflow	Cortex Workflow
Reporting Latency	6–12 hours	Minutes (near real-time)
Data Accuracy (Anomaly Detection)	70%	95%
IT Maintenance Cost (Annual)	\$1.2M	\$850K
Audit Preparation Time	3 weeks	5 days
Customer Satisfaction Index	72 / 100	88 / 100

4.7.1. NLP Capabilities in Snowflake Cortex

Snowflake Cortex integrates pre-trained large language models (LLMs) into the Snowflake Data Cloud, enabling companies to perform complex NLP tasks without requiring additional ML systems [21]. As a result of this seamless incorporation, latency issues, governance, and AI shareability are eliminated. Additionally, this tool has an innovative feature of automatically abstracting a long document, for instance, an agreement, claim description, and/or filing, into a miniature description. The feature is useful in sectors such as compliance and biotechnology, where representatives are tasked with analysing large volumes of text in a short period. As a model for document abstraction, transformer models (e.g., BERT, GPT) exploit semantic meaning instead of mere word frequency [22]. On a different topic, social media, consumer opinion, and research studies provide instant insights into consumer satisfaction and brand awareness. Text classification models (e.g., positive, neutral, or negative sentiment) for opinion mining, besides being generally applicable, can also specialise in specific domains such as biotechnology or finance. Cortex finds structured information such as names, dates, financial numbers, or medical classifications in unstructured text. It helps with analytics, fraud analysis, and compliance analysis. Entity recognition models use word embeddings and context-based models that enhance identification in challenging domains [22]. Unlike phrase-based searching, semantic searching is based on meaning. Cortex uses vector embeddings to map text into a high-dimensional vector space for context-based similarity searches. It helps a user discover relevant documents even when the actual phrases used are not known [23]. Cortex translates natural language queries into optimised SQL queries. It enables non-technical personnel to access data directly. It depends on the integrated intent recognition and query translation models in Snowflake [21].

4.8. Business Use Cases of Cortex

The inclusion of Natural Language Processing (NLP) capabilities in Snowflake Cortex has enabled a wide range of applications for businesses across industries. Snowflake Data Cloud has enabled Cortex to directly incorporate large language models,

enabling intelligent interaction with structured and unstructured data. Cortex assists in categorising support messages, identifying key issues, and providing automated responses [24]. The system reduces solve times while increasing consumer satisfaction by facilitating real-time evaluation and support prioritisation (Figure 7).



Figure 7: Business use cases of Cortex

4.8.1. Regulatory Compliance Monitoring

The contract analysis system uses various NLP models. These models help the system scan contracts, filed documents, or communications that may pose compliance risks. In this case, healthcare providers or financial institutions can automatically detect PHI leakage or flag clauses that do not comply with regulatory standards [25].

4.8.2. Financial Forecasting and Market Intelligence

Cortex can process text data from earnings releases, research analysis reports, and market news. Incorporating qualitative inputs into forecasting models helps increase the precision of market trend predictions and investment risk analysis [24].

4.8.3. Healthcare Claim and Provider Data Analysis

NLP functions evaluate provider notes, claim descriptions, and EDI/HL7 files to identify inconsistencies, fraudulent activity, or eligibility disputes. Automated extraction of medical codes in real time improves processing accuracy and enhances denial prevention. This technology facilitates payers to reduce denials and increase credibility in their business [22].

4.8.4. Human Resources and Talent Management

Cortex can process resumes, job descriptions, and employee feedback to pinpoint deficiencies in workforce skill sets. The inclusion of NLP processes in human resource management will support workforce planning and employee engagement [23].

4.8.5. Retail and E-Commerce Personalisation

Based on product reviews, queries, and browsing activities, it helps recommend products, optimise inventory, and increase customer engagement. This boosts customer loyalty and simplifies business operations [24]. Cortex simplifies case file analysis, identifies clauses, and highlights risk clauses for contracts, therefore speeding up contract analysis and risk detection [25].

5. Conclusion

The predictive analytics provided by Cortex further enable proactive financial planning and resource allocation, optimising capital deployment based on anticipated market trends and operational demands. This strategic foresight, coupled with the system's ability to automate time-consuming tasks, significantly expedites financial statement preparation and enables real-

time monitoring of key performance indicators. Snowflake Cortex represents a pivotal shift in enterprise data strategy. By replacing legacy reporting tools, augmenting ETL processes, and transforming databases into intelligent systems, Cortex enables organisations to build truly modern data ecosystems. As enterprises continue to navigate the complexities of digital transformation, Cortex offers a unified, AI-powered platform that not only meets today's demands but also anticipates tomorrow's needs. Snowflake Cortex represents a transformative advancement in data transfer. By embedding AI into ETL and database workflows, enterprises can achieve faster, more reliable, and cost-effective data movement. Its adoption marks a shift from static, manual processes to intelligent, adaptive ecosystems that redefine the value of data to the business.

The advent of Snowflake Cortex signals a major paradigm shift for companies seeking to use reporting, ETL, and database solutions. By seamlessly integrating artificial intelligence and machine learning algorithms within the Snowflake Data Cloud, Snowflake Cortex overcomes the inefficiencies that plagued traditional technologies that relied on batch-oriented ETL processes, traditional reporting software, and separate databases. With Snowflake Cortex, companies are poised to operate in an artificial intelligence-native world where natural language querying, automated SQL generation, semantic search, anomaly detection, and predictive analytics are built into their operations. In reporting, Cortex establishes a paradigm shift away from dashboards and toward predictive intelligence. In ETL, Cortex features an adaptive automation process that handles schema evolution and data quality checks without human interaction. At the database level, Cortex applies a paradigm shift by transforming passive storage into an intelligence engine, enabling self-optimisation and workload forecasting, as well as governance across every level of the data lifecycle.

Acknowledgement: The author gratefully acknowledges Harrisburg University of Science and Technology for providing an inspiring academic environment and the resources that supported this work. Sincere appreciation is extended to the faculty and staff for their guidance and encouragement throughout the research process.

Data Availability Statement: The data underlying the results of this study can be accessed by contacting the author.

Funding Statement: The author declares that this research was conducted without external financial assistance.

Conflicts of Interest Statement: There are no known conflicts of interest associated with this publication.

Ethics and Consent Statement: The author ensures ethical compliance and consents to the dissemination of this work for learning and research purposes.

References

1. S. Banerjee, "Advanced Data Management: A Comparative Study of Legacy ETL Systems and Unified Platforms," *Int. Res. J. Modernization in Eng. Tech. and Science*, vol. 6, no. 11, pp. 5677-5688, 2024.
2. M. B. Abughazala and H. Muccini, "Modelling Data Analytics Architecture for Smart Cities Data-Driven Applications using DAT," *arXiv:2307.08870v2*, 2023, Available: <https://arxiv.org/pdf/2307.08870v2> [Accessed by 10/10/2024].
3. H. A. Mohna, T. Barua, M. Mohiuddin, and M. M. Rahman, "AI-Ready Data Engineering Pipelines: A Review of Medallion Architecture," *Am. J. Scholarly Res. and Innovation*, vol. 1, no. 1, pp. 319-350, 2022.
4. S. C. Seethala, "Transforming Healthcare Data Warehouses with AI," *Int. J. Sci. Res. in Comp. Sci. Eng. and IT*, vol. 9, no. 2, pp. 749-752, 2023.
5. C. Althati, M. Tomar, and L. Shanmugam, "Enhancing Data Integration and Management: The Role of AI and ML in Modern Data Platforms," *Journal of Artificial Intelligence General Science JAIGS*, vol. 2, no. 1, pp. 220-233, 2024.
6. S. K. Nanjappan, "Machine Learning in Data Analytics and Reporting," *Int. J. for Res. in Applied Sci. and Eng. Tech.*, vol. 12, no. 9, pp. 1193-1200, 2024.
7. B. Khan, W. Khan, S. Jan, and M. I. Chughtai, "An Overview of ETL Techniques, Tools, Processes and Evaluations in Data Warehousing," *Journal on Big Data*, vol. 6, no. 1, pp. 1-20, 2024.
8. S. Gangarapu and V. V. R. Chilukoori, "The Future of Data Warehousing: Trends and Challenges," *Int. J. Sci. Res. in Comp. Sci. Eng. and IT*, vol. 10, no. 5, pp. 470-479, 2024.
9. R. H. Chowdhury, "Blockchain and AI: Driving the Future of Data Security and BI," *World J. of Advanced Research and Reviews*, vol. 23, no. 1, pp. 2559-2570, 2024.
10. S. K. Singu, "A Comprehensive Approach to ML Integration in Data Warehousing," *Journal of Technology and Systems*, vol. 6, no. 6, pp. 28-37, 2024.
11. S. Dashora, N. Gupta, M. P. S. Tanwar, A. Singh, and S. Kaur, "Cloud-based Data Analytics for Business Intelligence," *Int. J. for Res. in Applied Sci. and Eng. Tech.*, vol. 11, no. 11, pp. 2758-2765, 2023.
12. G. B. Komarina and N. K. R. Chopra, "SAP Analytics Cloud – Transforming BI with AI and ML," *World J. of Advanced Research and Reviews*, vol. 26, no. 2, pp. 3256-3262, 2025.

13. E. O. Eboigbe, O. Farayola, F. O. Olatoye, N. O. Chinwe, and C. Daraojimba, "Business Intelligence Transformation Through AI and Data Analytics," *Eng. Sci. and Tech. J.*, vol. 4, no. 5, pp. 285-307, 2023.
14. M. R. Sokkula and S. K. Vuppala, "Implementing Machine Learning for ETL Data Transformation and Anomaly Detection," *Int. J. for Multidisciplinary Research*, vol. 6, no. 6, pp. 45-52, 2024.
15. A. Mbata, Y. Sripada, and M. Zhong, "A Survey of Pipeline Tools for Data Engineering," *arXiv:2406.08335*, 2024, Available: <https://arxiv.org/abs/2406.08335> [Accessed by 12/10/2024].
16. V. Ananthkrishnan, S. S. Kunju, and R. Pachyappan, "Autonomous GenAI Agents for Legacy-to-Cloud ETL Modernization," *Journal of Artificial Intelligence General Science*, vol. 1, no. 1, pp. 274-290, 2024.
17. P. S. Dhoni, "From Data to Decisions: Enhancing Retail with AI and ML," *Int. J. of Computing and Engineering*, vol. 5, no. 1, pp. 38-51, 2024.
18. M. Wilbrod, I. Alex, and K. Matovu, "ICT Adoption and Business Performance: A Case Study of Retailers in Kampala," vol. 3, no. 10, pp. 560-580, 2024.
19. C. P. Efunniyi, A. O. Abhulimen, A. N. Obiki-Osafiele, O. S. Osundare, E. E. Agu, and I. A. Adeniran, "Strengthening Corporate Governance and Financial Compliance," *Finance and Accounting Research Journal*, vol. 6, no. 8, pp. 1597-1616, 2024.
20. B. O. Antwi, B. O. Adelokun, and A. O. Eziefule, "Transforming Financial Reporting with AI," *Int. J. of Advanced Economics*, vol. 6, no. 6, pp. 205-223, 2024.
21. Snowflake Inc., "Snowflake Cortex AI," *Snowflake Blog*, 2025, Available: <https://www.snowflake.com/en/product/features/cortex/> [Accessed by 07/10/2025].
22. J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," *arXiv:1810.04805*, 2019, Available: <https://arxiv.org/abs/1810.04805> [Accessed by 24/10/2024].
23. T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient Estimation of Word Representations in Vector Space," *arXiv:1301.3781*, 2013, Available: <https://arxiv.org/abs/1301.3781> [Accessed by 07/10/2024].
24. A. McAfee and E. Brynjolfsson, "Big Data: The Management Revolution," *Harvard Business Review*, vol. 90, no. 10, pp. 60-68, 2012.
25. S. Ponnusamy and P. Gupta, "Connecting the Dots: How Data Lineage Helps in Effective Data Governance," *International Journal on Computer Science and Engineering*, vol. 10, no. 10, pp. 6-10, 2023.