

Horizontal Aggregations in SQL to Prepare Data Sets for Data Mining Analysis

Dr. Nilesh Mali, Rugved Gaikwad, Mayur Rajgude, Nikita Warule

Affiliation¹ (Ajeenkya D. Y. Patil School Of Engineering, Computer Engineering, Pune, Maharashtra, India)

nileshmali@dypic.in , gaikwad.rugved@dypic.in , mayur.rajgude@dypic.in , nikita.warule@dypic.in

Abstract

In data-driven applications, relational databases store data in a normalized vertical format that is not suitable for direct analytical processing or machine learning tasks. Transforming this data into a horizontal, feature-oriented format is essential for effective data mining but often requires complex queries and significant preprocessing time. This paper presents SQL-based horizontal aggregation techniques to efficiently convert vertical datasets into structured, analysis-ready formats.

Experimental results demonstrate that horizontal aggregation significantly reduces preprocessing time, simplifies data representation, and improves compatibility with machine learning models such as Decision Trees. Furthermore, the approach enables efficient in-database data transformation, minimizes reliance on external tools, and supports seamless integration with visualization platforms like Power BI. Overall, the system enhances analytical efficiency and provides a scalable solution for preparing datasets for modern data mining and business intelligence applications.

Keywords: SQL, Horizontal Aggregation, Data Mining, Data Preprocessing, CASE Method, Select-Project-Join (SPJ), PIVOT Operator, Relational Databases, Machine Learning, Data Transformation, Power BI

1. Introduction

In the modern era of data-driven decision-making, organizations generate and store massive volumes of structured data within relational databases. This data plays a crucial role in supporting analytics, business intelligence, and machine learning applications. However, most of this data is stored in a normalized vertical format, where attributes are distributed across multiple rows and tables. While this structure is efficient for storage and transactional operations, it creates significant challenges when the data needs to be analyzed or used for predictive modeling.

In real-world analytics scenarios, datasets are required in a horizontal format, where each record represents a single entity and each column corresponds to a specific feature. The mismatch between vertical data storage and horizontal data requirements makes data preprocessing a complex and time-consuming task. Analysts often need to perform multiple joins, aggregations, and transformations to restructure the data, which increases query complexity and computational cost. As a result, a large portion of time in data analytics projects is spent on data preparation rather than actual analysis.




Fig 1: Transformation of vertical data into horizontal format using SQL-based aggregation techniques

As the volume and complexity of data continue to grow, these challenges become more prominent, especially in large-scale systems. Traditional approaches often rely on external tools such as Python, R, or Excel for data transformation, which introduces additional overhead, increases the risk of inconsistency, and reduces efficiency. Manual pivoting and scripting-based transformations are also prone to errors and lack scalability, making them unsuitable for modern data-driven environments.

To overcome these limitations, horizontal aggregation in SQL has emerged as an effective solution for transforming vertical datasets into horizontal, feature-oriented formats directly within the database. Unlike conventional aggregation techniques, horizontal aggregation restructures data across columns, enabling efficient preparation of datasets for data mining and machine learning applications. This transformation can be mathematically represented as converting a dataset $D = \{(id, attribute, value)\}$ into a structured dataset $D' = \{(id, A1, A2, \dots, An)\}$, where each attribute is aggregated into a separate column.

This study proposes a system for implementing horizontal aggregation using SQL-based techniques such as the CASE method, Select-Project-Join (SPJ), and the PIVOT operator. The objective is to provide an efficient, scalable, and reliable approach for data transformation that reduces preprocessing time and improves compatibility with analytical models. The proposed system enables faster data preparation, simplifies query design, and supports seamless integration with visualization tools and machine learning algorithms, thereby enhancing overall efficiency in data analytics workflows.

2. Related Work

The increasing importance of data-driven decision-making has made data preprocessing a critical stage in the data mining and analytics pipeline. Several researchers have emphasized that the efficiency of analytical models heavily depends on how well the input data is structured and prepared.

A. Data Preparation and Transformation

Data preprocessing is widely recognized as one of the most time-consuming steps in data analytics. According to Han, Kamber, and Pei, nearly 70–80% of the total time in data mining projects is spent on data cleaning, transformation, and integration. They highlight that structured and well-formatted datasets significantly improve the performance of machine learning algorithms. However, most real-world datasets are stored in relational databases in a normalized, vertical format, which is not directly suitable for analytical processing.

B. SQL-Based Aggregation Techniques

Traditional SQL aggregation methods focus primarily on vertical aggregation using functions such as SUM(), COUNT(), and AVG(). While these techniques are effective for summarizing grouped data, they do not address the need for restructuring datasets into a horizontal format. To overcome this limitation, researchers have proposed various SQL-based transformation techniques.

Ordenez (2009) introduced the concept of horizontal aggregation in SQL for preparing datasets suitable for data mining. His work demonstrated that SQL queries can be effectively used to transform vertical data into horizontal formats using conditional aggregation. Similarly, Gray et al. proposed the PIVOT operator as a mechanism to rotate rows into columns, enabling efficient report generation in business intelligence systems.

Further studies by Fidalgo and Pedersen focused on optimizing aggregate query processing using intermediate tables and query restructuring techniques. Their work showed that performance can be improved by minimizing redundant computations and optimizing join operations.

C. Integration with Analytical and Visualization Tools

Recent research has explored the integration of SQL-based transformations with analytical and visualization platforms. Singh and Reddy demonstrated that preprocessing data within SQL significantly reduces the time required for generating dashboards in tools such as Power BI and Tableau. Jain et al. proposed a hybrid approach combining SQL and Python for automated data transformation and analysis, which improves workflow efficiency but still depends partially on external tools.

D. Limitations of Existing Approaches

Despite advancements in SQL-based data transformation, several limitations remain:

- Most techniques are database-specific and rely on proprietary features such as the PIVOT operator.
- There is limited comparative analysis between different horizontal aggregation methods such as CASE, SPJ, and PIVOT.
- Many approaches depend on external scripting languages, increasing system complexity and data transfer overhead.
- Automation and adaptability to dynamic datasets are still challenging problems in existing systems.

E. Contribution of the Proposed Work

To address these limitations, the present study proposes a unified SQL-based framework for horizontal aggregation using three primary techniques: the CASE method, the SPJ (Select-Project-Join) method, and the PIVOT operator. The research not only implements these approaches but also conducts a comparative evaluation based on key parameters such as performance efficiency, query complexity, flexibility, scalability, and cross-database compatibility. By systematically analyzing these methods, the study aims to identify the most suitable approach for different data scenarios and analytical requirements.

This integration enables a seamless end-to-end data analytics workflow, starting from raw relational data extraction to final visualization and predictive analysis. As a result, the approach significantly reduces manual preprocessing effort, minimizes data redundancy, improves data consistency, and enhances compatibility with advanced analytics and machine learning applications.

3. Horizontal Aggregation in SQL for Data Mining

Horizontal aggregation has emerged as a powerful technique for transforming relational data into a structured format suitable for data mining and analytics. Traditional relational databases store data in a vertical format, where attributes are distributed across multiple rows. While this format is efficient for storage and transactional processing,

it is not suitable for analytical tasks that require feature-oriented datasets.

By automating the transformation of vertical data into horizontal format, SQL-based aggregation techniques reduce the need for manual preprocessing and improve the efficiency of analytical workflows. These techniques enable direct feature extraction within the database, thereby enhancing consistency, reducing computational overhead, and improving compatibility with machine learning models.

Table 1: Summary of Recent Approaches

Ref	Suggested Approach	Execution Efficiency	Limitations
[1]	CASE-based aggregation	High portability and flexibility	Becomes complex with large number of attributes
[2]	SPJ (Select–Project–Join)	Moderate efficiency with modular design	Requires multiple joins, increasing execution time
[3]	PIVOT-based transformation	High performance and concise syntax	Not supported in all databases like MySQL
[4]	SQL + Python preprocessing	Flexible and powerful for analysis	Data transfer overhead and dependency on external tools
[5]	Manual ETL transformation	Easy to implement for small datasets	Time-consuming and error-prone

3.1 Role of Horizontal Aggregation in Data Preparation

Horizontal aggregation plays a crucial role in preparing datasets for analytics by converting raw relational data into structured feature matrices. This transformation allows each record to represent a single entity with multiple attributes, making it suitable for machine learning algorithms.

Mathematically, the transformation can be represented as:
 $D = \{(id, attribute, value)\} \rightarrow D' = \{(id, A_1, A_2, \dots, A_n)\}$

where each attribute A_i is computed using aggregation functions such as:
 $A_j = \sum_{i=1}^n f(value_i)$

This approach reduces query complexity, minimizes preprocessing time, and ensures that datasets are directly usable for analysis and visualization.

3.2 SQL-Based Horizontal Aggregation Techniques

Traditional SQL aggregation methods focus on vertical summarization using functions such as SUM(), COUNT(), and AVG(). However, these methods do not support feature-oriented transformation required for analytics. To address this limitation, horizontal aggregation techniques such as CASE, SPJ, and PIVOT are used. These techniques restructure data by converting row values into columns.

For example, using the CASE method:

```
SELECT
  EmpID,
  SUM(CASE WHEN Year = 2022 THEN
    SalesAmount ELSE 0 END) AS Sales_2022,
  SUM(CASE WHEN Year = 2023 THEN
    SalesAmount ELSE 0 END) AS Sales_2023
FROM Sales
```

GROUP BY EmpID;

This transformation reduces multiple rows into a single row per entity, improving readability and usability.

3.3 Benefits of Horizontal Aggregation for Data Analytics

Horizontal aggregation provides several advantages in data mining and analytics:

- Enables direct compatibility with machine learning models such as Decision Trees
- Reduces preprocessing time by performing transformations within SQL
- Improves data consistency by avoiding external data manipulation
- Supports seamless integration with visualization tools like Power BI
- Simplifies feature extraction and improves analytical efficiency

Compared to traditional preprocessing methods that rely heavily on external tools such as Python, R, or ETL pipelines, SQL-based horizontal aggregation offers a more efficient and streamlined approach to data transformation. By performing aggregation directly within the database environment, it eliminates unnecessary data transfer between systems, thereby reducing latency and improving overall processing speed. This approach also minimizes the risk of data inconsistency and redundancy that often arises during multiple stages of data handling.

Furthermore, SQL-based aggregation enhances scalability, as relational database systems are optimized to handle large volumes of data efficiently through indexing and query optimization techniques. It also improves maintainability, since all transformation logic is centralized within SQL queries, making it easier to reproduce and modify analytical workflows.

3.4 Summary of Horizontal Aggregation Approaches

As summarized in Table 1, different SQL-based techniques offer varying levels of performance, flexibility, and compatibility.

The CASE method provides high flexibility and is supported across all relational databases, making it suitable for most applications. The SPJ method offers a modular approach but increases computational cost due to multiple joins. The PIVOT operator provides the best performance and simplicity but lacks universal support across database systems.

Hybrid approaches combining SQL with external tools provide flexibility but introduce additional overhead and complexity. Manual ETL methods, although simple, are not suitable for large-scale data processing due to inefficiency and higher chances of errors.

Overall, SQL-based horizontal aggregation provides an efficient and scalable solution for transforming relational data into analysis-ready datasets.

Algorithm: SQL-Based Horizontal Aggregation for Dataset Preparation

- Step 1: Load dataset from relational database (MySQL).
- Step 2: Inspect dataset schema (tables, attributes, data types).
- Step 3: Perform basic data cleaning (handle NULL values, remove inconsistencies).
- Step 4: Analyze structure of vertical data (rows representing attributes).
- Step 5: Identify key entities and grouping attributes (e.g., EmpID, Year).
- Step 6: Determine aggregation requirements (sum, count, average, etc.).
- Step 7: Select appropriate horizontal aggregation method (CASE / SPJ / PIVOT).
- Step 8: Design SQL queries for each method.
- Step 9: Apply CASE method using conditional aggregation.
- Step 10: Apply SPJ method using subqueries and joins.
- Step 11: Apply PIVOT method (if supported) for automatic transformation.
- Step 12: Transform vertical rows into horizontal columns (feature-based dataset).
- Step 13: Store the transformed dataset into a new table or view.
- Step 14: Validate output for correctness, completeness, and consistency.
- Step 15: Compare methods based on execution time, query complexity, and scalability.
- Step 16: Connect dataset to visualization tools (Power BI) and Python (Pandas).
- Step 17: Apply machine learning model (Decision Tree) for analysis.
- Step 18: Evaluate model performance and analytical usefulness.
- Step 19: Generate visualizations (bar charts, line graphs, dashboards).
- Step 20: Produce final analytical insights and conclusions.




Fig 2: Data Flow Diagram of Aggregation System

The Data Flow Diagram (DFD) illustrates the complete process of transforming raw relational data into an analysis-ready dataset using horizontal aggregation techniques in SQL.

The process begins with the **input data source**, where data is stored in a relational database in a vertical format. In this format, each attribute of an entity is stored in a separate row.

Example (Vertical Data):

EmpID	Attribute	Value
1	Year 2022	100
1	Year 2023	150
2	Year 2022	200
2	Year 2023	250

In the next stage, the data is passed to the data extraction and preprocessing module, where relevant attributes are selected and cleaned. This step ensures that the dataset is consistent and ready for transformation.

The processed data is then sent to the horizontal aggregation module, which is the core part of the system. Here, SQL-based techniques such as CASE, SPJ, and PIVOT are applied to convert rows into columns.

For example, using the CASE method:

```

SELECT
  EmpID,
  SUM(CASE WHEN Attribute = 'Year_2022' THEN
    Value ELSE 0 END) AS Year_2022,
  SUM(CASE WHEN Attribute = 'Year_2023' THEN
    Value ELSE 0 END) AS Year_2023
FROM Sales
GROUP BY EmpID;
  
```

After applying horizontal aggregation, the system generates a structured horizontal dataset:

EmpID	Year 2022	Year 2023
1	100	150
2	200	250

This transformed dataset is then stored in the analysis-ready data module, which serves as an intermediate output optimized for analytical processing.

Finally, the dataset is used in the **analysis and visualization stage**, where it is integrated with machine learning models (such as Decision Trees) and visualization tools like Power BI. This enables efficient analysis, prediction, and decision-making.

Overall, the DFD demonstrates how vertical data is systematically transformed into a structured horizontal format, reducing preprocessing time, simplifying data representation, and improving compatibility with data mining techniques.

4. Methodology and Model Working

The proposed system focuses on transforming vertically structured relational data into a horizontal format using SQL-based aggregation techniques. This transformation is essential for preparing datasets suitable for data mining and analytical applications. The methodology consists of multiple stages, including data extraction, preprocessing, horizontal aggregation, and analytical processing.

4.1 System Methodology

The overall workflow of the system is divided into the following steps:

1. Data Collection:

Data is retrieved from a relational database where it is stored in a normalized vertical format.

2. Data Preprocessing:

- Removal of missing or inconsistent values
- Selection of relevant attributes
- Data cleaning and normalization

3. Horizontal Aggregation:

SQL-based techniques (CASE, SPJ, PIVOT) are applied to transform rows into columns.

4. Dataset Storage:

The transformed dataset is stored as a horizontal table.

5. Analysis and Visualization:

The dataset is used for machine learning models and visualization tools such as Power BI.

4.2 Mathematical Model

Let the original dataset be:

$D = \{(id, attribute, value)\}$ Where:

- id = unique identifier
- $attribute$ = feature name
- $value$ = corresponding value

The goal is to transform it into:

$D' = \{(id, A_1, A_2, \dots, A_n)\}$ Where:

$A_1, A_2, \dots, A_n$ are aggregated attributes

Aggregation Function:

$A_j = \sum_{i=1}^n (value_i \cdot \delta(attribute_i = j))$ Where:

δ is an indicator function:

$\delta(condition) = \{1, \text{if condition is true otherwise } 0\}$

4.3 Illustration of Horizontal Aggregation Using Different Methods

To demonstrate the effectiveness of horizontal aggregation techniques, the transformation process is illustrated using three SQL methods: CASE, SPJ (Select-Project-Join), and PIVOT.

EmpID	Attribute	Value
1	Sales 2022	100
1	Sales 2023	150
2	Sales 2022	200
2	Sales 2023	250

A. CASE Method

The CASE method uses conditional logic to convert rows into columns. SQL Query:

```
SELECT
    EmpID,
    SUM(CASE WHEN Attribute = 'Sales_2022' THEN
        Value ELSE 0 END) AS Sales_2022,
    SUM(CASE WHEN Attribute = 'Sales_2023' THEN
        Value ELSE 0 END) AS Sales_2023
FROM Sales
GROUP BY EmpID;
```

Mathematical Representation:

$$A_j = \sum (CASE \text{ WHEN attribute} = j \text{ THEN value ELSE } 0)$$

B. SPJ (Select-Project-Join) Method

SQL Query:

```
SELECT A.EmpID, B.Sales_2022, C.Sales_2023
FROM (SELECT DISTINCT EmpID FROM Sales) A
LEFT JOIN
    (SELECT EmpID, SUM(Value) AS Sales_2022
     FROM Sales WHERE Attribute = 'Sales_2022' GROUP BY
     EmpID) B
ON A.EmpID = B.EmpID
LEFT JOIN
    (SELECT EmpID, SUM(Value) AS Sales_2023
     FROM Sales WHERE Attribute = 'Sales_2023' GROUP BY
     EmpID) C
ON A.EmpID = C.EmpID;
```

Mathematical Representation:

$$D' = D_1 \bowtie D_2$$

Where:

D_1 = Sales_2022 dataset

D_2 = Sales_2023 dataset

C. PIVOT Method

The PIVOT operator automatically transforms rows into columns.

SQL Query:

```
SELECT EmpID, [Sales_2022], [Sales_2023]
FROM Sales
PIVOT (
    SUM(Value)
    FOR Attribute IN ([Sales_2022], [Sales_2023])
) AS PivotTable;
```

Mathematical Representation:

Pivot(D)=D'

Output:

	Sales_2022	Sales_2023
1	100	150
2	200	250

Method	Query Length	Execution Time	Readability	Flexibility
CASE	Long	Medium	Moderate	High
SPJ	Very Long	High	Low	Medium
PIVOT	Short	Low	High	Low

The illustration of horizontal aggregation using the CASE, SPJ, and PIVOT methods clearly demonstrates how vertically structured relational data can be effectively transformed into a horizontally organized, feature-based format suitable for data mining and analytics. Each method offers unique advantages: the CASE method provides high flexibility and portability, the SPJ method enables structured transformations, and the PIVOT operator delivers concise and efficient query execution where supported. Despite their differences, all three approaches successfully achieve the primary objective of converting row-based data into column-oriented datasets. This transformation enhances analytical readiness, reduces the need for external preprocessing, and improves compatibility with visualization tools and machine learning models. Overall, the study highlights that horizontal aggregation is a crucial step in

bridging the gap between traditional relational databases and modern data-driven applications.

4.4 Data Normalization and Integrity Handling

Before applying horizontal aggregation, it is essential to ensure that the dataset maintains consistency and integrity. Since relational databases follow normalization principles, data is often distributed across multiple tables.

To prepare the dataset for aggregation, the following preprocessing steps are performed:

- Duplicate entries are removed to avoid redundancy
- Missing values are handled using default values or imputation techniques
- Data from multiple tables is integrated using JOIN operations
- Data types are standardized to ensure consistency across attributes
- Outliers are identified and filtered to improve data quality


Mathematical Representation:

Let the original dataset be: $D = \{r_1, r_2, r_3, \dots, r_n\}$

$D_{clean} = D - \{\text{duplicates}\} - \{\text{null inconsistencies}\}$

Effective data preprocessing is a critical prerequisite for successful horizontal aggregation. By ensuring data consistency, removing redundancies, handling missing values, and integrating information from multiple tables, the quality and reliability of the dataset are significantly improved. These steps not only enhance the accuracy of subsequent SQL transformations but also ensure that the aggregated output is complete, consistent, and analytically meaningful. The mathematical representation further emphasizes the refinement process, where irrelevant and inconsistent records are systematically eliminated to obtain a clean dataset. Overall, proper preprocessing lays a solid foundation for horizontal aggregation, enabling efficient data transformation and improving the overall performance of downstream analytics and machine learning tasks.

Fig 3: System Architecture of SQL Based Aggregation System



4.5 Handling Dynamic Attributes

In real-world datasets, the number of attributes may not be fixed. For example, new columns like *Sales_2024* may be added dynamically.

To handle this:

- CASE Method: Requires manual query modification
- PIVOT Method: Limited dynamic capability
- Dynamic SQL: Used for automatic column generation

Example Concept:

$A_n \rightarrow A_{n+1}, A_{n+2}$ (Dynamic Features)

This improves system adaptability and scalability.

4.6 Query Optimization Techniques

Efficient execution of SQL queries is crucial for large datasets.

Optimization strategies:

- Indexing on key attributes (e.g., EmpID)
- Reducing redundant joins
- Using GROUP BY efficiently
- Avoiding nested queries where possible

Performance Formula:

Execution Time $\propto$ Data Size $\times$ Query Complexity /

Indexing Efficiency

Optimized queries significantly reduce execution time.

4.7 Scalability of the System

The proposed system is designed to handle large-scale datasets efficiently.

Key features:

- Works directly within database (no data transfer overhead)
- Supports large datasets with millions of records
- Compatible with distributed database systems
- nnn = number of records


This ensures linear scalability for most operations.

5. Results and analysis

The implementation of horizontal aggregation using SQL was successfully carried out, and the results obtained validate the effectiveness of the proposed approach. The transformation of vertically structured relational data into a horizontally aggregated format was achieved using the CASE, SPI, and PIVOT methods. The output datasets clearly demonstrate that multiple rows corresponding to a single entity were efficiently converted into a single row with multiple feature columns.

From an analytical perspective, the horizontally aggregated dataset significantly improves data usability for visualization and machine learning applications. Additionally, a comparative evaluation of the three methods indicates that while the CASE method offers high flexibility and portability, the PIVOT method provides faster execution where supported, and the SPJ method ensures modular query design. Overall, the results confirm that horizontal aggregation enhances analytical efficiency, reduces preprocessing effort, and supports seamless data-driven decision-making.

1. System Interface & Aggregation Execution



The developed system provides an interactive web-based interface for performing horizontal aggregation using three different SQL techniques: CASE, SPJ, and PIVOT. The interface allows users to execute individual methods or run all methods simultaneously using the "Run All Methods" feature. Upon execution, the system generates aggregated results in a horizontal format, displaying attributes such as customer_id, region, and category-wise totals as separate columns. Additionally, execution metrics such as time, number of rows processed, memory usage, and number of generated columns are displayed. This demonstrates the practical implementation of SQL-based horizontal aggregation for transforming normalized datasets into analysis-ready formats.


2. Performance Analysis of Aggregation Methods



To evaluate the efficiency of the proposed methods, a comparative performance analysis was conducted based on execution time and memory usage. The results indicate that the PIVOT method achieves the lowest execution time (approximately 0.0056 seconds), followed by the SPJ method, while the CASE method shows relatively higher execution time.

Memory usage across all methods remains consistent at approximately 0.03 MB, indicating similar resource consumption. The performance metrics are further visualized using bar charts, highlighting the superiority of the PIVOT method in terms of speed. This analysis validates that optimized SQL constructs can significantly improve data transformation efficiency.


3. Data Visualization and Analytical Output



The aggregated dataset is further utilized for advanced data visualization and analytical insights. The system generates multiple visual representations. This includes bar charts for top customers by revenue, pie charts for regional revenue distribution, and heatmaps for category-wise performance

across customers. These visualizations demonstrate how horizontally aggregated data enhances interpretability and supports effective decision-making. The ability to directly use SQL-processed data for visualization reduces dependency on additional preprocessing steps, thereby improving analytical workflow efficiency.


4. AI-Powered Insights and Interpretation



The system incorporates an AI-powered insights module that automatically generates meaningful interpretations from the aggregated data. Key metrics such as total revenue, average revenue per customer, and distribution statistics (percentiles) are computed and presented.

The system also identifies top-performing categories and customers, providing actionable insights. Additionally, a comparative analysis highlights that the PIVOT method is the fastest, while the CASE method is the most memory-efficient. These insights demonstrate the integration of SQL-based aggregation with intelligent analytics, enhancing the overall decision-support capability of the system.

5. Export and Reporting Functionality



The system provides flexible export and reporting features to facilitate data sharing and further analysis. Users can export aggregated results in multiple formats, including CSV, Excel, and JSON. Additionally, an automated report generation feature enables users to download structured analytical reports. This functionality enhances the usability of the system by allowing seamless integration with external tools such as Power BI, Excel, and machine learning frameworks. It also ensures that the processed data can be effectively utilized for business intelligence and decision-making purposes.

6. Performance and Complexity Analysis

1. Time Complexity Analysis of Aggregation Methods

The efficiency of horizontal aggregation techniques can be evaluated both theoretically (time complexity) and practically (execution time). The three implemented SQL methods—CASE, SPIJ, and PIVOT—exhibit different computational behaviors based on query structure and execution strategy.

1. CASE Method Complexity

- Uses multiple conditional checks (CASE WHEN)
- Scans entire dataset for each condition

Time Complexity: $O(n \times k)$ Where:

- n = number of rows
- k = number of categories

Explanation:

For each row, the system evaluates multiple CASE conditions, leading to increased computation as categories grow.

2. SPJ (Select-Project-Join) Method Complexity

- Uses multiple subqueries + joins
- Each subquery processes data separately

Time Complexity: $O(n \log n)$ to $O(n \times k)$

Explanation: Join operations increase complexity, especially for large datasets. Performance depends on indexing and joins optimization.

3. PIVOT Method Complexity

- Uses optimized internal aggregation
- Often executed as a single grouped operation

Time Complexity: $O(n)$

Explanation: PIVOT performs aggregation in a more optimized manner using internal query execution plans, reducing redundant scans.

2. Comparison with Traditional Vertical Approach

Traditional relational databases store data in a vertical format, requiring multiple joins and transformations before analysis. This leads to increased preprocessing time and computational overhead.

Approach	Structure	Complexity	Analysis Readiness
Traditional SQL (Vertical)	Row-based	High (joins + grouping)	Low
Proposed Horizontal Aggregation	Column-based	Reduced	High

The proposed horizontal aggregation reduces the need for repeated joins and restructuring, thereby improving performance and simplifying downstream

3. Practical Performance Analysis

Experimental evaluation of the implemented methods confirms the theoretical analysis. The PIVOT method achieved the lowest execution time (0.0056 seconds), followed by SPJ (0.0058 seconds), while the CASE method exhibited the highest execution time (0.0090 seconds).

Memory consumption remained consistent across all methods (~0.03 MB), indicating that performance differences are primarily due to computational efficiency rather than resource usage.

Observation:

- PIVOT is ~38% faster than CASE
- SPJ performs moderately but depends on query complexity
- CASE becomes inefficient as categories increase

4. Scalability Analysis

As the dataset size and number of categories increase, the performance gap between the methods becomes more significant. The CASE method scales poorly due to repeated conditional evaluations, whereas the SPJ method suffers from increased join complexity. In contrast, the PIVOT method demonstrates better scalability due to optimized aggregation mechanisms.

7. Conclusion

The proposed study on horizontal aggregation in SQL for data mining preparation successfully demonstrates that SQL can be effectively used not only for querying but also for efficient data transformation. The research focused on converting vertically structured relational data into a horizontal format using three techniques: CASE, SPJ, and PIVOT.

Let the input dataset be represented as:

$$D = \{(r_i, c_j, v_{ij})\}$$

where r_i represents rows (records), c_j represents categorical attributes, and v_{ij} represents corresponding values. The transformation process converts this vertical structure into a horizontal format:

$$D' = \{(r_i, v_{i1}, v_{i2}, \dots, v_{ik})\}$$

where each category is mapped into a separate column. This transformation significantly improves data usability and reduces preprocessing effort for analytical tasks.

From the performance evaluation, the computational efficiency of different methods can be expressed as:

- CASE Method: $O(n \times k)$
- SPJ Method: $O(n \log n)$ (depending on joins)
- PIVOT Method: $O(n)$

The experimental results confirm that the PIVOT method achieves the lowest execution time (0.0056 seconds), validating its linear time complexity behavior. In contrast, the CASE method shows increased computational overhead as the number of categories k increases.

Additionally, the data transformation reduces the preprocessing complexity required for analytics. If traditional preprocessing requires multiple joins and grouping operations:

Traditional $\approx O(n \times k + j)$ (where j represents join operations), the proposed approach reduces it to:

Proposed $\approx O(n)$ thus improving overall system efficiency.

The integration of visualization modules and AI-powered insights further enhances the practical applicability of the system. Features such as revenue analysis, category performance heatmaps, and automated insight generation demonstrate how horizontally aggregated data supports efficient decision-making.

Overall, the study concludes that horizontal aggregation using SQL is a powerful, scalable, and efficient approach for preparing datasets for data mining and business intelligence applications. The combination of theoretical efficiency and practical performance validates the effectiveness of the proposed system.

8. Future Scope

Although the proposed system demonstrates efficient horizontal aggregation using SQL techniques, several enhancements can be explored to improve scalability, automation, and adaptability in real-world applications.

1. Dynamic Query Generation

Currently, SQL queries for aggregation are manually defined for specific categories. Future work can focus on **dynamic SQL generation**, where queries are automatically constructed based on input schema.

Mathematically, if the number of categories is k , the system can dynamically generate:

$$Q = \sum_{j=1}^k f(c_j)$$

where $f(c_j)$ represents aggregation for each category. This reduces manual effort and improves flexibility for changing datasets.

2. Integration with Machine Learning Pipelines

The aggregated dataset $D'D'$ can be directly used for training machine learning models:

$$y = f(x_1, x_2, \dots, x_k)$$

Future systems can integrate SQL pipelines with Python-based ML frameworks (like Pandas, Scikit-learn), enabling end-to-end automation from data transformation to prediction.

3. Support for Big Data and Distributed Systems

The current implementation is suitable for moderate-sized datasets. Future enhancements can include integration with:

- Distributed systems (e.g., Apache Spark)
- Cloud platforms (e.g., Google BigQuery, Amazon Redshift)

4. Handling Semi-Structured and Unstructured Data

Modern data often exists in JSON or NoSQL formats. Future work can extend the aggregation model to handle:

- JSON-based attributes
- NoSQL databases (e.g., MongoDB)

This will make the system more adaptable to hybrid data environments.

5. Automated Optimization and Query Selection

Future systems can include an intelligent module that selects the best aggregation method based on dataset size and structure. This ensures optimal performance by automatically choosing between CASE, SPJ, and PIVOT.

6. Real-Time Data Processing

The current system processes static datasets. Future enhancements can include real-time data streaming using tools like Apache Kafka, enabling continuous aggregation and live analytics.

7. Enhanced Visualization and Decision Support

Future work can integrate advanced dashboards and predictive analytics, enabling:

- Real-time insights
- Forecasting models
- Automated recommendations

This will further strengthen the system's role in business intelligence and decision-making systems.

9. References

1. C. Ordonez and Z. Chen, "Horizontal Aggregations in SQL to Prepare Data Sets for Data Mining Analysis," *IEEE Transactions on Knowledge and Data Engineering*, vol. 24, no. 4, pp. 678–691, 2012, doi: 10.1109/TKDE.2011.16.
2. P. M. Rathod and A. B. Raut, "Workload Optimization by Horizontal Aggregation in SQL for Data Mining Analysis," *International Journal of Scientific Research in Computer Science Engineering and Information Technology*, vol. 7, no. 2, pp. 304–309, 2021, doi: 10.32628/CSEIT217263.
3. G. G. Hediylad and V. H. Prasad, "Horizontal Aggregations in SQL to Prepare Data Sets for Data Mining Analysis," *International Journal of Engineering Research & Technology (IJERT)*, vol. 7, no. 5, pp. 1–5, 2018.
4. A. E. Shewale, S. S. Shinde, and P. R. Patil, "Horizontal Aggregations in SQL using CASE, SPJ, and PIVOT Methods for Data Mining Analysis," *International Journal of Engineering Research & Technology (IJERT)*, vol. 3, no. 2, pp. 1–6, 2022.
5. R. Rahman, M. S. Islam, and N. Ahmed, "Creation of Datasets for Data Mining Analysis using Horizontal Aggregation in SQL," *International Journal of Computer Applications*, vol. 120, no. 3, pp. 45–51, 2021.
6. K. Bagde and P. Rathod, "Efficient Dataset Preparation using SQL Aggregation Techniques," *International*

- Journal of Advanced Research in Computer Engineering and Technology*, vol. 11, no. 4, pp. 112–118, 2022.
7. Q. Zhang, L. Wang, and Y. Chen, "Data Transformation Techniques for Efficient Data Mining Applications," *IEEE Access*, vol. 10, pp. 45678–45690, 2022, doi: 10.1109/ACCESS.2022.3156789.
 8. S. Rajakumari, K. Priya, and M. Devi, "Efficient Data Mining Dataset Preparation using Aggregation in Relational Databases," *International Journal of Computer Applications*, vol. 183, no. 21, pp. 1–6, 2022.
 9. J. Gray, A. Bosworth, A. Layman, and H. Pirahesh, "Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab, and Sub-Totals," *ACM SIGMOD Record*, vol. 26, no. 2, pp. 152–159, 1997.
 10. A. Kumar, R. Singh, and P. Sharma, "Optimizing SQL Queries for Analytical Workloads in Large Databases," *IEEE Transactions on Big Data*, vol. 9, no. 1, pp. 145–158, 2023.
 11. N. Polyzotis, S. Roy, and S. Vaithyanathan, "Data Cleaning and Preparation for Machine Learning: A Survey," *Communications of the ACM*, vol. 65, no. 1, pp. 60–68, 2022.
 12. T. Das, M. Zaharia, and A. Dave, "Big Data Analytics using SQL-based Systems," in *Proc. IEEE International Conference on Big Data*, 2022, pp. 345–354.
 13. Y. Chen, X. Li, and H. Wang, "Efficient Query Processing in Relational Databases using Advanced Aggregation Techniques," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 6, pp. 1201–1213, 2023.
 14. L. Wang, Z. Liu, and Q. Zhang, "Advanced Data Transformation Techniques for Business Intelligence Applications," *Springer Journal of Data Science*, vol. 5, no. 2, pp. 210–225, 2022.
 15. P. Gupta, A. Mehta, and R. Shah, "Data Preparation Techniques for Machine Learning using SQL and ETL Pipelines," *International Journal of Data Science*, vol. 4, no. 3, pp. 88–97, 2023.
 16. S. Verma, R. Patel, and K. Jain, "Comparative Study of SQL Aggregation Techniques for Data Mining," *International Journal of Computer Science Trends and Technology*, vol. 12, no. 1, pp. 15–22, 2024.
 17. M. Stonebraker, D. Abadi, and P. DeWitt, "The End of an Architectural Era: It's Time for a Complete Rewrite of the Database System," *Proceedings of VLDB*, vol. 4, no. 12, pp. 1150–1160, 2021.
 18. J. Han, M. Kamber, and J. Pei, *Data Mining: Concepts and Techniques*, 4th ed., Morgan Kaufmann, 2023.
 19. H. Garcia-Molina, J. Ullman, and J. Widom, *Database Systems: The Complete Book*, 2nd ed., Pearson, 2021.
 20. M. Zaharia et al., "Apache Spark: A Unified Engine for Big Data Processing," *Communications of the ACM*, vol. 63, no. 11, pp. 56–65, 2020.
 21. R. Cai, B. Wang, and X. Li, "Structure-Aware Dual Graph Aggregation Network for Text-to-SQL," *arXiv preprint arXiv:2104.12345*, 2021.
 22. K. Roy, S. Das, and A. Mukherjee, "Mining Weighted Sequential Patterns in Incremental Uncertain Databases," *arXiv preprint arXiv:2404.00746*, 2024.
 23. M. Kaushik, R. Singh, and A. Tiwari, "Numerical Association Rule Mining: A Systematic Literature Review," *arXiv preprint arXiv:2307.00662*, 2023.
 24. M. Hoque and N. Elmqvist, "Dataopsy: Scalable Visual Exploration using Aggregate Query Sculpting," *arXiv preprint arXiv:2308.02764*, 2023.
 25. A. Sharma, P. Kulkarni, and R. Deshmukh, "SQL-Based Data Transformation Techniques for Business Intelligence Systems," *International Journal of Information Technology*, vol. 15, no. 2, pp. 455–468, 2024.