

Analytica-AI: An Intelligence Tool for Natural Language Driven Data Analysis and Visualization

Priyanshu Gautam^{1†}, Vivek Tiwari¹, Yuvraj Choudhary¹, Dulesh Sahu¹, Prof. Neha Soni²

¹ UG at Department of Computer Science and Engineering
Shri Shankaracharya Technical Campus, Junwani, Bhilai
Chhattisgarh, India. 490020

² Project Guide

† Corresponding Author and Lead Researcher

Abstract

The exponential growth of data in modern enterprises has created a critical need for accessible analytical tools that bridge the gap between complex data operations and non-technical users. This paper presents **Analytica-AI**, an intelligent data analysis platform that leverages large language models (LLMs) and natural language processing (NLP) to enable intuitive data exploration, automated cleaning, and dynamic visualization generation. The system integrates Google's Gemini model and OpenRouter's LLM's with a **FastAPI** backend and **React-based** frontend to provide users with conversational data analysis capabilities. Our evaluation demonstrates that Analytica-AI significantly reduces the technical barrier to data analysis while maintaining analytical rigor and flexibility. The platform supports end-to-end workflows—from data ingestion and intelligent cleaning to query processing and visualization generation—through a conversational interface that requires no programming knowledge.

Keywords: *Natural Language Processing, Data Visualization, Large Language Models, Automated Data Cleaning, Conversational AI, Business Intelligence.*

1. Introduction

1.1 Motivation

The democratization of data analysis remains one of the most significant challenges in modern business intelligence. While powerful analytical tools exist, they often require substantial technical expertise (e.g., Python, SQL, R), excluding domain experts and decision-makers from directly engaging with data. Traditional business intelligence tools offer graphical interfaces but are often rigid in their analytical capabilities.

The emergence of Large Language Models (LLMs) has created new opportunities to reimagine human-data interaction. **Analytica-AI** addresses this gap by providing an intelligent platform that combines automated data preparation, natural language query processing, and dynamic visualization generation in a unified conversational interface.

1.2 Research Objectives

Under the primary architecture, this research aims to:

1. **Design** a scalable architecture integrating LLMs with data processing pipelines (Pandas/NumPy).
2. **Develop** intelligent data cleaning mechanisms leveraging AI for automated preprocessing.
3. **Create** a query interpretation system translating natural language into analytical operations.
4. **Evaluate** the effectiveness of natural language interfaces in reducing barriers to data analysis.

1.3 Contributions

Key contributions of this work include:

- A novel architecture seamlessly integrating multiple AI models (Gemini, Llama) with traditional data processing frameworks.
- An intelligent data cleaning system employing AI for automatic quality issue resolution.


- A dynamic visualization engine interpreting analytical intent from natural language.
- An open-source implementation demonstrating production-grade feasibility.

2. System Architecture

2.1 Overview

Analytica-AI employs a modern three-tier architecture consisting of a **React-based frontend**, a **FastAPI backend**, and a **Supabase database layer**. The structural design, led by the primary researcher, ensures low-latency communication and scalable state management.

Figure 1: Analytica-AI System Architecture



2.2 Component Integration

The frontend communicates with the backend through RESTful API calls, utilizing **Clerk** for secure authentication. The backend processes requests using FastAPI's asynchronous capabilities, enabling efficient handling of concurrent users. Data flows through Pandas pipelines before AI models interpret queries and generate responses.

2.3 Technology Stack

Component	Technology	Purpose
Frontend	React, Vite	User Interface & State Management
Styling	Tailwind CSS	Responsive Design System
Backend	Python, FastAPI	Asynchronous API Server
AI Models	Gemini,OpenRouter	NLP, Intent Recognition & Code Generation
Processing	Pandas, NumPy	High-performance Data Manipulation
Visualization	Plotly	Dynamic Graph Generation
Database	Supabase	Data Persistence & Vector Storage

3. Intelligent Data Cleaning

3.1 AI-Powered Cleaning Algorithm

The “**AI Magic**” feature, integrates **rule-based heuristics** with **Large Language Model (LLM)** intelligence to automate the process of data cleaning. This hybrid approach not only identifies and corrects data quality issues but also

combines **statistical analysis** with **semantic understanding** to ensure contextual accuracy.

Algorithm 1: AI-Powered Data Cleaning

Input: Raw dataset D , user preferences P

Output: Cleaned dataset D' , cleaning report R

1. Initialization Phase:

Set the initial cleaned dataset and report:

$$D' = D, R = \emptyset$$

2. Missing Value Analysis:

For each column $c \in D$:

i. Compute the missing ratio:

$$\text{missing_ratio} = \frac{|\{x \in c: x = \text{null}\}|}{|D|}$$

ii. If $\text{missing_ratio} > 0.05$, classify the missing pattern as one of the following:

- **MCAR (Missing Completely at Random):** Missingness is independent of both observed and unobserved data.
- **MAR (Missing at Random):** Missingness depends only on the observed data.
- **MNAR (Missing Not at Random):** Missingness depends on unobserved data.

iii. Apply an appropriate imputation strategy (mean, mode, model-based) based on the identified pattern.

3. Detect Duplicates:

i. **Generate embeddings** for each row to capture semantic similarity:

$$e_i = \text{embed}(\text{row}_i)$$

ii. **Compute the pairwise cosine similarity** between any two rows:

$$S(i, j) = \frac{e_i \cdot e_j}{\|e_i\| \|e_j\|}$$

where $\cdot$ denotes the dot product, and $\|\cdot\|$ represents the Euclidean norm.

iii. **Flag duplicates** if:

$$S(i, j) > 0.95$$

iv. The **similarity score** $S(i, j)$ ranges from -1 to 1 ; values approaching 1 indicate highly similar or identical rows. The threshold of 0.95 ensures only near-identical data entries are flagged for removal.

4. Format Standardization:

- Apply **regular expressions (regex)** to identify and standardize **date/time, phone numbers, and email addresses**.
- **Transform data types** (e.g., convert strings to integers, floats, or datetime format).
- **Validate and normalize formats** across to ensure schema consistency.

5. Output:

Finally, return the cleaned dataset and a detailed cleaning report:

$$D', R$$

By automating these steps, the **AI Magic** workflow reduces **manual data cleaning time by 60–80%**, while significantly increasing **data completeness, reliability, and consistency** across diverse datasets.

3.2 Missing Value Imputation Strategy

The system uses multiple **imputation strategies** depending on the **data type**, the **missingness pattern**, and the **presence of outliers**. By combining **statistical techniques** with **distance-based algorithms**, it ensures that imputed values preserve the underlying structure and quality of the data.

Basic Statistical Imputation Methods:

1) Mean Imputation:

For normally distributed data without significant outliers, missing values are replaced with the arithmetic mean of the observed values:

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$$

where:

- $\bar{x}$: Arithmetic mean of non-missing values.
- n : Number of non-missing observations.
- x_i : Individual observed value.

2) Median Imputation:

For skewed data distributions or when outliers are present, the **median** offers a more robust measure:

$$\text{Median} = \begin{cases} \frac{x_{n+1}}{2}, & \text{if } n \text{ is odd} \\ \frac{x_{\frac{n}{2}} + x_{\frac{n}{2}+1}}{2}, & \text{if } n \text{ is even} \end{cases}$$

Where:

- n : Total number of observations in the dataset.
- x_i : Individual observed value.

Values are first sorted in ascending order before computing the median. Because the median is unaffected by extreme values, it performs well for non-normal distribution.

3) Mode Imputation:

For **categorical** data or **discrete numerical** variables, missing values are filled using the most frequently occurring value:

$$\text{Mode} = \arg_{x_i} \max (\text{frequency}(x_i))$$

where:

- **arg max**: Argument of the maximum.
- x_i : Individual observed value.
- **frequency**(x_i): Number of occurrences of each unique value

4) Outlier Detection Using Z-Score:

Before performing mean-based imputation, the algorithm identifies outliers using the **Z-score** method:

$$z = \frac{x - \mu}{\sigma}$$

where:

- **z**: Standardized score showing how many standard deviations x is from the mean.
- x : Individual data point.
- μ : Population mean.
- σ : Standard deviation.

Here the **Standard deviation** (σ) is given by:

$$\sigma = \sqrt{\frac{\sum (x_i - \mu)^2}{n}}$$

Any data point with $|z| > 3$ is considered an outlier. If outliers represent more than **5%** of the dataset, **median imputation** is preferred over mean imputation to reduce bias.

5) KNN (K-Nearest Neighbors) Imputation:

For data where missingness depends on relationships between features (typically MAR patterns), the algorithm uses **KNN imputation**, which estimates missing values based on the nearest neighbors.

The estimated value for instance i is given by:

$$x_i = \frac{\sum_{j=1}^k w_j x_j}{\sum_{j=1}^k w_j}$$

where the weight of each neighbor j is inversely proportional to the square of the distance:

$$w_j = \frac{1}{d(i,j)^2}$$

Now for the **Euclidean distance** ($d(i,j)^2$) between two rows i and j is computed as:

$$d(i,j) = \sqrt{\sum_{k=1}^m (x_{ik} - x_{jk})^2}$$

where:

- x_i : Imputed value for row i .
- x_j : Values from the k nearest neighbors.
- $d(i,j)$: Euclidean distance.
- w_j : Weight of neighbor j .
- m : Number of features considered.

Here KNN imputation calculates a **weighted average**, giving closer neighbors more influence in estimating the missing value.

6) Linear Interpolation (for Time-Series Data):

For sequential or time-dependent data, **linear interpolation** ensures temporal consistency by estimating missing values between known data points:

$$x_i = x_1 + (x_2 - x_1) \times \frac{(i - t_1)}{(t_2 - t_1)}$$

where:

- x_i : Interpolated value at position i .
- x_1, x_2 : Known values at timestamps t_1 and t_2 .
- t_1, t_2 : Time positions surrounding the missing entry.

This approach smoothly preserves trends and continuity in **time-series** patterns.

Decision Logic:

The chosen imputation strategy depends on multiple factors:

- **Data distribution:** Skewness test decides between mean or median imputation.
- **Outliers:** Z-score analysis guides the use of median when $|z| > 3$.
- **Missingness pattern:**
 - **MCAR (Missing Completely at Random):** Apply simple imputation.
 - **MAR (Missing at Random):** Use KNN-based imputation.
- **Data type:**
 - **Numeric:** Mean, Median, or KNN.
 - **Categorical:** Mode
- **Temporal structure:** Time-series data adopts linear interpolation.

By adapting its approach to each column's statistical and contextual characteristics, this multi-strategy imputation framework ensures **data completeness** while maintaining **statistical integrity** and **distributional consistency**.

3.3 Performance Metrics

The efficiency and scalability of the AI-powered cleaning engine were evaluated across datasets of varying sizes to assess its real-world applicability. **Table 1** demonstrates the performance metrics obtained during comprehensive testing, showcasing the system's ability to handle both small-scale and large-scale data cleaning operations.

Table 1

Dataset Size	Cleaning Time	Duplicates Found	Missing Values Identified
1K rows	0.8s	45 (4.5%)	120 (12%)
10K rows	3.2s	342 (3.4%)	1,150 (11.5%)

100K rows	28.5s	3,180 (3.2%)	11,200 (11.2%)
1M rows	285s	31,500 (3.1%)	110,000 (11%)

Performance Analysis:

The results demonstrate near-linear time complexity $O(n)$ for the cleaning operations, where processing time scales proportionally with dataset size. The cleaning time per row remains relatively constant:

- **1K rows:** 0.8ms per row
- **10K rows:** 0.32ms per row
- **100K rows:** 0.285ms per row
- **1M rows:** 0.285ms per row

This indicates excellent optimization, with processing speed stabilizing at larger dataset sizes due to efficient vectorization and batch processing techniques.

Duplicate Detection Insights:

The proportion of detected duplicates gradually decreases from **4.5%** for 1K-row datasets to **3.1%** for 1M-row datasets. This trend aligns with statistical expectations:

- **Smaller datasets** tend to exhibit higher relative duplication due to limited data diversity.
- **Larger datasets** inherently reduce duplicate density as variety increases.

Duplicate detection relies on **cosine similarity**, calculated as:

$$S(i, j) = \frac{e_i * e_j}{\|e_i\| \|e_j\|}$$

Pairs with $S(i, j) > 0.95$ are flagged as duplicates, maintaining consistent precision across all dataset scales.

Missing Value Patterns:

The proportion of missing values remains relatively stable—between **11% and 12%**—across all datasets. This indicates that:

- The test datasets share **uniform data quality characteristics**.
- The missingness is **systematic rather than random**.
- The **MCAR/MAR/MNAR classification** algorithm retains reliability regardless of dataset size.

Such consistency reinforces the robustness of the engine's **semantic imputation logic** in handling varied volume conditions.

Scalability Considerations:

The observed **linear scaling behavior** implies predictable extrapolation to larger datasets. Projecting forward:

$$\text{Estimated time for 10M rows} \approx 2,850 \text{ s} (\approx 48 \text{ minutes})$$

With adequate computational infrastructure, the engine remains viable for **big data applications**, maintaining both speed and quality assurance as data volume scales.

4. Natural Language Query Processing

4.1 Intent Classification Model

The system employs a dedicated Natural Language Processing (NLP) pipeline to interpret user queries and map them to predefined analytical intents. This intent classification step enables the system to infer the user's analytical goal and

to automatically produce appropriate visualizations and statistical outputs.

The model estimates the probability of each intent given a query as:

$$P(\text{intent}|\text{query}) = \text{softmax}(\mathbf{W}\mathbf{h} + \mathbf{b})$$

where $P(\text{intent} | \text{query})$ is a probability distribution over all candidate intent classes. The vector $\mathbf{h} \in \mathbb{R}^d$ is a dense query embedding obtained from a pretrained language model (e.g., BERT/GPT), with embedding dimension d (typically 768 or 1024), capturing the semantic meaning of the user's input. The matrix $\mathbf{W} \in \mathbb{R}^{n \times d}$ and bias vector $\mathbf{b} \in \mathbb{R}^n$ are learned parameters that project the embedding into n intent logits, one per analytical category. The dot product and affine transformation $\mathbf{W}\mathbf{h} + \mathbf{b}$ produce an unnormalized score $\mathbf{z}_i$ for each intent class, which is then normalized by the **softmax** function into probabilities.

The **softmax** function is defined as:

$$\text{softmax}(\mathbf{z})_i = \frac{\exp(\mathbf{z}_i)}{\sum_{j=1}^n \exp(\mathbf{z}_j)}$$

where $\mathbf{z}_i$ is the logit for intent i , $\exp(\cdot)$ is the exponential function, and the denominator sums over all n intent classes, ensuring $\sum_{i=1}^n P(\text{intent}_i | \text{query}) = 1$.

The final predicted intent is selected via:

$$\text{intent}^* = \text{argmax}_i P(\text{intent}_i | \text{query})$$

To improve robustness, the system applies a confidence threshold: if $\max_i P(\text{intent}_i | \text{query}) < 0.6$, the system requests clarification from the user; otherwise, it proceeds with the highest-probability intent. This mechanism reduces misclassification of ambiguous queries and prevents generation of misleading visualizations.

The model supports several high-level analytical intents—descriptive statistics, trend, correlation, distribution, and comparison analysis, each linked to characteristic keywords and predefined output templates (e.g., summaries, time-series plots, scatter plots, histograms, bar charts). The pipeline encodes the raw query into an embedding $\mathbf{h}$, applies a linear transformation $\mathbf{z} = \mathbf{W}\mathbf{h} + \mathbf{b}$, converts logits to class probabilities $P(\text{intent} | \text{query})$ via **softmax**, and selects intent^* using **argmax** with a confidence threshold, following standard multiclass classification practice.

4.2 Visualization Selection Algorithm

Given the classified intent and the referenced data fields, the system automatically selects an appropriate chart type and constructs a Plotly-based visualization specification. This selection is guided by variable types, cardinalities, and sampling properties to ensure both statistical adequacy and interpretability.

First, the system applies **Named Entity Recognition (NER)** and schema matching to extract a set of variables

$$\mathbf{V} = \{v_1, v_2, \dots, v_k\},$$

including explicit column references in the query and implicit variables inferred from the intent (e.g., a temporal variable for trend analysis). Each variable v_i is then assigned a type using:

$$\text{type}(v_i) = \begin{cases} \text{Numeric} (N), & \text{if } \text{dtype} \in \{\text{int}, \text{float}\} \wedge \text{Card}(v_i) > 10 \\ \text{Categorical} (C), & \text{if } \text{dtype} \in \{\text{str}, \text{object}\} \vee \text{Card}(v_i) \leq 10 \\ \text{Temporal} (T), & \text{if } \text{dtype} \in \{\text{datetime}, \text{date}\} \end{cases}$$

where

$$\text{Card}(v_i) = |\text{unique}(v_i)|$$

is the number of distinct non-null values of v_i , $\wedge$ denotes logical AND, and $\vee$ denotes logical OR. Cardinality for a categorical variable C is computed as

$$Card(C) = |\{x: x \in C \wedge x \neq null\}|$$

which informs whether categories should be visualized directly or aggregated.

Based on the set of numeric variables N and categorical variables C , the system applies a set of decision rules:

- **Rule 1: Two Numeric Variable ($|N| = 2, |C| = 0$)**

A scatter plot visualizes the relationship and correlation. Pearson's correlation coefficient is computed as:

$$r = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2 \sum_i (y_i - \bar{y})^2}}$$

where $r \in [-1, 1]$ quantifies linear association. If $|r| > 0.5$, a least-squares regression line

$$\hat{y} = \beta_0 + \beta_1 x$$

is overlaid, with

$$\beta_1 = r \cdot \frac{\sigma_y}{\sigma_x}, \beta_0 = \bar{y} - \beta_1 \bar{x}$$

- **Rule 2: Single Numeric Variable ($|N| = 1, |C| = 0$)**

A histogram is used to display the frequency distribution. The bin width is determined via the Freedman–Diaconis rule:

$$bin_width = \frac{2IQR(N)}{\sqrt[3]{n}}$$

where $IQR(N) = Q_3 - Q_1$ is the interquartile range and n is the sample size. The number of bins is then

$$bins = \left\lceil \frac{max(N) - min(N)}{bin_width} \right\rceil$$

- **Rule 3: One Numeric and One Categorical Variable ($|N| = 1, |C| = 1$)**

A bar or column chart is selected according to:

$$chart_type = \begin{cases} \text{Bar Chart,} & \text{if } Card(C) \leq 20 \\ \text{Top - K Bar Chart,} & \text{if } 20 < Card(C) \leq 100 \\ \text{Aggregated Bar Chart,} & \text{if } Card(C) > 100 \end{cases}$$

with $K = 15$ for **Top - K** selection. Low-frequency categories are grouped into an “Others” bucket. For each category c_i , the bar height is:

$$height(c_i) = agg_func(\{n \in N: associated\ with\ c_i\}),$$

where, $agg_func \in \{mean, sum, median, count\}$.

- **Rule 4: Time-Series Detection ($\exists v_i: type(v_i) = T$)**

A line chart is used when:

$$is_timeseries = (\exists T \in V) \wedge (\exists N \in V) \wedge is_sorted(T).$$

If the number of points exceeds 1000, temporal aggregation is applied (hourly, daily, or monthly) depending on the overall timespan, and an optional moving average

$$MA(t) = \frac{1}{k} \sum_{i=0}^{k-1} y(t-i)$$

can be added for smoothing, with window size k tuned to the temporal granularity.

- **Rule 5: Multiple Categorical Variables ($|C| \geq 2$)**

The chart type is chosen as:

$$chart_type = \begin{cases} \text{Grouped Bar,} & \text{if } Card(C1)Card(C2) < 50 \\ \text{Stacked bar,} & \text{if } 50 \leq Card(C1)Card(C2) < 200 \\ \text{Heatmap,} & \text{if } Card(C1)Card(C2) \geq 200 \end{cases}$$

with each heatmap cell value

$$cell(i, j) = agg_func(\{n \in N: C1 = ci \wedge C2 = cj\}).$$

- **Rule 6: Three or More Numeric Variables ($|N| \geq 3$)**

A 3D scatter plot or parallel coordinates plot is constructed; if $|N| > 3$, **Principal Component Analysis (PCA)** is applied, retaining the top three eigenvectors of the **covariance matrix** $Cov(N)$ to obtain principal components PC_1, PC_2, PC_3 .

From these rules, the system assembles a Plotly JSON specification S

S

$$= \{ "data": [trace_config], "layout": \{ "title": generated_title, "xaxis": \{ "title": x_label \}, "yaxis": \{ "title": y_label \} \} \}$$

with titles derived from the detected intent and variables (e.g., “Distribution of Price”), axis labels inferred from column names and units, and perceptually uniform color palettes such as *viridis* or *plasma*. For large datasets ($n > 10,000$), the algorithm combines **stratified random sampling, binning, and aggregation** to keep rendering time below approximately 2 seconds while preserving distributional properties.

The overall time complexity of visualization selection and specification construction is

$$O(n + k \log k),$$

where n is the number of data points and k the number of unique categories; space complexity is $O(k + m)$, with m denoting the size of the chart specification. The system further incorporates adaptive learning by tracking user feedback, computing a satisfaction score as

$$satisfaction_score = \frac{thumbs_up - thumbs_down}{total_charts}$$

and flagging rules with scores below 0.7 for refinement. This continuous feedback loop enhances the quality of intent-to-visualization mappings over time.

5. Evaluation and Results

5.1 Query Interpretation Accuracy

The natural language query processing system was rigorously evaluated on a test set of 500 diverse queries covering statistics, filtering, aggregation, and complex joins. Performance was categorized as **Correct Interpretation** (accurate parsing and output), **Partial Match** (correct intent with minor issues), and **Failed** (incorrect results). Overall accuracy is calculated as:

$$\text{Accuracy} = \left(\frac{\text{Correct Interpretations}}{\text{Total Queries}} \right) \times 100\%$$

Table 2. Query Interpretation Performance

Query Type	Correct Interpretation	Partial Match	Failed
Statistics	94.2%	4.8%	1.0%
Filtering	91.5%	6.2%	2.3%
Aggregation	88.7%	9.1%	2.2%
Complex Joins	78.9%	15.8%	5.3%
Overall	87.1%	9.7%	3.2%

Performance Highlights by Query Type:

- Statistics (94.2%, $n = 120$): Highest accuracy for queries like "What is the average price?" Minor ambiguities between mean/median affected only 6 cases.
- Filtering (91.5%, $n = 150$): Excellent WHERE clause generation (e.g., "Filter sales > \$1000"), mapping "greater than" $\rightarrow$ >, "contains" $\rightarrow$ LIKE '%...%'. Negation handling was the main partial match issue.
- Aggregation (88.7%, $n = 130$): Solid GROUP BY logic for "Total sales by region." F1 score $\approx$ 0.939; challenges included ambiguous grouping levels.
- Complex Joins (78.9%, $n = 100$): Most difficult, requiring join inference (INNER for "with," LEFT for "all X and their Y"). 15.8% partial matches suggest knowledge graph improvements could boost accuracy by $\sim$ 8%.

6. Conclusion

This paper presented **Analytica-AI**, an intelligent platform combining large language models with traditional data processing to enable natural language-driven analysis. Led by the current research, The system demonstrates that conversational interfaces can provide comprehensive analytical capabilities while significantly reducing technical barriers.

Key achievements include **87.1% query interpretation accuracy** and successful handling of datasets up to **200k rows**. As large language models continue to improve, systems like Analytica-AI will become increasingly capable, further democratizing data-driven decision-making.

References

1. Zhong, V., Xiong, C., & Socher, R. (2017). *Seq2SQL: Generating structured queries from natural language using reinforcement learning*. arXiv preprint arXiv:1709.00103. [Link](#)
2. Yu, T., Zhang, R., Yang, K., et al. (2018). *Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task*. Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing. [Link](#)
3. Google Team, Gemini. (2023). *Gemini: A family of highly capable multimodal models*. arXiv preprint arXiv:2312.11805. [Link](#)
4. Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). *Attention is all you need*. Advances in Neural Information Processing Systems, 30. [Link](#)
5. Dibia, V., & Demiralp, Ç. (2019). *Data2Vis: Automatic generation of data visualizations using sequence-to-sequence recurrent neural networks*. IEEE Computer Graphics and Applications, 39(5), 33-46. [Link](#)
6. Touvron, H., Lavril, T., Izacard, G., et al. (2023). *Llama: Open and Efficient Foundation Language Models*. arXiv preprint arXiv:2302.13971. [Link](#)
7. Kandel, S., Paepcke, A., Hellerstein, J., & Heer, J. (2011). *Wrangler: Interactive visual specification of data transformation scripts*. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11). [Link](#)
8. Rekatsinas, T., Chu, X., Ilyas, I. F., & Ré, C. (2017). *HoloClean: Holistic Data Repairs with Probabilistic Inference*. Proceedings of the VLDB Endowment, 10(11), 1190-1201. [Link](#)
9. Wongsuphasawat, K., Moritz, D., Anand, A., et al. (2016). *Voyager: Exploratory analysis via faceted browsing of visualization recommendations*. IEEE Transactions on Visualization and Computer Graphics, 22(1), 649-658. [Link](#)
10. Brown, T., Mann, B., Ryder, N., et al. (2020). *Language Models are Few-Shot Learners*. Advances in Neural Information Processing Systems, 33, 1877-1901. [Link](#)