



Analyst Buddy AI: A Review of AI-Driven Techniques for Automated Data Analytics, Visualization and Insight Generation

Kamakshi Hanvat¹, Gauri Yadav², Suhani Yadav³, and Nakshatra Sharma⁴, Dr. U. P. Akare, Assistant Professor⁵

⁵Under the Guidance of

Department of Artificial Intelligence & Data Science Engineering, Priyadarshini College of Engineering, Nagpur (R.T.M. Nagpur University), Maharashtra, India

ABSTRACT :

Turning an uploaded spreadsheet into something a decision-maker can act on still takes several separate tools and a fair amount of technical skill. This review looks at ten studies published between 2015 and 2025 that cover visualization recommendation, natural-language interfaces, LLM-based analysis agents, interactive dashboards, data-quality tooling, automated machine learning, customer segmentation and AI-supported business intelligence. We summarise what each study contributes, compare them in one table, and relate them to Analyst Buddy AI, a Flask-based platform that carries CSV and Excel files through validation, preprocessing, exploratory analysis, KPI dashboards, natural-language questions, K-Means segmentation, sales forecasting and PDF reporting. Across the literature we find five recurring gaps: work that stops at a single stage of the analytics pipeline, unreliable or unexplained language-model output, limited attention to data quality before modelling, weak and inconsistent evaluation of clustering and prediction choices, and thin treatment of privacy and usability. We end with a research agenda for a unified, testable and privacy-aware analytics assistant.

Keywords— *Data Analytics, Visualization Recommendation, Natural-Language Interfaces, Large Language Models, Data Quality, Customer Segmentation, K-Means Clustering, Business Intelligence.*

1. Introduction

Most organisations hold more data than they can comfortably interpret. A typical analysis starts with a raw file and passes through inspection, cleaning, exploration, chart selection, KPI calculation, modelling and finally a written report. Each hand-off between tools costs time, and each one is a place where a non-specialist can get stuck.

The literature has attacked these steps one at a time. Visualization recommenders try to choose charts automatically [1], and recent work asks language models to do the same from a plain-English question [2]. Dashboard research looks at how findings should be presented to different audiences [11]. Text-to-SQL and agent-based systems aim to let people query databases and even draft analysis reports without writing code [4], [10]. Separate strands study data quality [6], automated model design [5], customer segmentation [9], [8] and the wider role of AI in business intelligence [14].

Two observations motivate a closer look. The first is a skills gap. Xu et al. note that analysing relational data usually calls for SQL knowledge and trained data scientists, expertise that many micro, small and medium-sized companies cannot afford [10]. The same barrier faces anyone who simply has a spreadsheet and a question. Business intelligence has long been limited by how much data people can process and interpret by hand, and AI techniques are widely expected to widen that capacity [14]. The second observation is that automation can fail quietly. Data quality affects the performance, fairness and robustness of machine-learning models [6], and even a well-studied task such as customer segmentation still leaves open questions about which algorithm to use, how many segments to create and how to judge the result [9], [8].

We should be clear about the scope of this review. The ten studies were not written with an upload-and-report platform in mind. They were selected because each one deals with a stage that Analyst Buddy AI also covers, and they range from a 2015 vision paper on visualization recommendation to work published in 2025 on LLM-based analysis agents. We therefore treat them as design evidence: ideas worth borrowing, results that set realistic expectations, and pitfalls to avoid. The contributions of this review are threefold. It summarises the ten studies by theme, it compares them in a single table that links each one to a module of the platform, and it draws out five gaps that lead to a problem statement for future work.

This review brings those strands together and asks what they mean for a single, lightweight platform. Analyst Buddy AI accepts CSV and Excel files and offers automated preprocessing, KPI and dashboard generation, natural-language question answering, K-Means customer segmentation, sales forecasting and report export. Section 2 reviews the ten studies by theme and links each to the platform. Section 3 states the gaps that remain. Section 4 concludes and outlines future work.

2. Literature Review

The ten studies are grouped below by the stage of the analytics workflow they address, from choosing charts to interpreting results for business users. Table 1 gives one row per study, with its technique, its main reported point and the part of Analyst Buddy AI it informs. The numbers quoted are those reported by the original authors; the studies use different datasets and metrics, so they are read side by side as design evidence and not ranked against each other.

Table 1. Single-table comparison of the reviewed studies and Analyst Buddy AI

Ref.	Study	Focus / technique	Key reported point	Relevance to Analyst Buddy AI
<i>Visualization recommendation and dashboards</i>				
[1]	Vartak et al. (2015)	Vision paper on visualization recommendation systems; SeeDB prototype	Manual chart building leaves large parts of a dataset unexplored. Recommendations depend on five axes and quality criteria such as surprise, and need pruning and sampling to stay interactive.	Basis for choosing charts from column types and user intent
[2], [12]	Sah et al. (2024)	GPT-4 prompt that turns a table and a natural-language query into an analytic specification	87.02% accuracy on 740 queries against 64.05% for NL4DV, but about 25 s per answer instead of 3 s. Errors included broken JSON and axis labels that did not match the plotted field.	Shows how to explain why a chart was suggested, and why axes must be checked
[11]	Grant, Trevino and Perdue (2025)	Three-phase dashboard workflow (planning, development, deployment) using Tableau and R Shiny	Asking the intended users what they need, and testing with them before release, is part of the process itself.	Guides dashboard design and usability testing
<i>Natural-language interfaces and analysis agents</i>				
[4], [13]	Liu et al. (2025)	Review of natural-language interfaces to databases (Text-to-SQL)	Work splits into understanding the question and building the query; rule-based, machine-learning, pre-trained and LLM methods are compared on accuracy, generalisation, expressiveness and cost.	Helps decide how questions become data operations
[10]	Xu et al. (2025), DAgent	LLM agent with planning, tool and memory modules for report generation	Breaks a question into sub-queries, tunes the SQL and writes a report; beat the tested baselines on completeness, correctness and conciseness.	Question decomposition and report writing; works on uploaded files, not databases
<i>Data quality and automated machine learning</i>				
[6]	Zhou et al. (2024)	Survey of 17 data-quality tools by dimension, metric and function	Standard exploratory checks are not enough for large data; proposes a roadmap for open-source tools and sees LLMs as a possible aid.	Supports the preprocessing and validation layer
[5], [7]	Baratchi et al. (2024)	Overview of AutoML: search space, search strategy, performance evaluation	Covers hyperparameter search, architecture search, existing systems and open challenges.	Points to automatic comparison of forecasting models
<i>Customer segmentation</i>				
[9]	Sulthana et al. (2024)	RFM features with K-Means and other clustering methods on two datasets	Silhouette score peaked at $k = 5$ (0.4234) while the elbow method suggested $k = 4$, so the metrics can disagree.	K-Means module and the choice of K

Ref.	Study	Focus / technique	Key reported point	Relevance to Analyst Buddy AI
[8], [3]	Salminen et al. (2023)	Systematic review of 172 articles on algorithmic segmentation	46 algorithms and 14 metrics found; K-means most common; experts seldom judged the results.	Justifies K-Means as a default and expert review of clusters
<i>AI in business intelligence</i>				
[14]	Rane et al. (2024)	Review of ML, NLP and big data in business intelligence	Highlights privacy, consent, anonymisation and fairness concerns alongside predictive and real-time analytics.	Frames the target users and the privacy requirement
<i>This work</i>				
—	Analyst Buddy AI	Flask, pandas, Plotly, K-Means, forecasting, generative-AI query layer	Prototype; evaluation plan only, no results yet.	Upload-to-report analytics in one workflow

2.1 Analyst Buddy AI: Applying the Patterns

Analyst Buddy AI is built with a Flask backend, pandas for data handling and Plotly for interactive charts. A dataset moves through one pipeline:

User → Upload CSV/Excel → Validation → Preprocessing → EDA → KPI and Dashboard → AI/NLP Query → ML Modules → PDF Report

The guiding rule is that numbers come from code, not from the language model. Aggregation, filtering and clustering run in programmatic components, and the model is used to interpret and explain. The machine-learning layer offers K-Means segmentation and sales forecasting, and the reporting layer combines KPIs, charts, model output and written insights into a downloadable PDF. The architecture is modular, so each stage can be tested and improved on its own.

2.2 Cross-Cutting Synthesis

Each of the ten studies concentrates on one part of the analytics workflow: choosing charts [1], [2], presenting dashboards [11], querying data [4], [10], cleaning it [6], modelling it [5], [7], [8], [9], or setting the strategic context [14]. Evaluation is just as varied. Reference [2] measures accuracy on a query corpus, [10] uses a purpose-built benchmark scored on completeness, correctness and conciseness, [9] relies on clustering indices, [8] codes the literature, and [1] and [14] are conceptual. Table 1 puts the studies and Analyst Buddy AI side by side.

3. Problem Statement

Reading the ten studies together points to five gaps, and each one affects a platform such as Analyst Buddy AI.

- **Fragmentation across the pipeline.** Recommenders [1], [2], dashboards [11], query systems [4], [10], quality tools [6], AutoML [5] and segmentation studies [9], [8] each address one stage. Even the most complete system, DAgent [10], works on relational databases rather than uploaded files, and stops at report generation.
- **Reliability and explainability of language-model output.** The best NL-to-chart result in the corpus is 87.02% accurate [2], and its errors include mislabeled axes and invalid specifications. Correctness is also a named evaluation criterion in [10]. A wrong number in an analytics answer is costly, so every figure should come from an actual computation.
- **Limited attention to data quality in end-to-end tools.** Reference [6] treats quality as central, yet the visualization and querying papers largely assume clean input. Automated cleaning in an end-to-end platform needs to be visible and reversible.
- **Weak handling of model and cluster choices.** Choosing the number of segments is often reduced to one hyperparameter judged by one metric [8], and the evidence in [9] shows that different measures can disagree. Domain experts rarely join the evaluation [8].
- **Privacy and usability treated lightly.** Privacy and ethics are discussed at review level [14], while dashboard research stresses user input and testing [11]. Neither is reflected in most technical systems.

Problem statement. The field has strong components but few lightweight systems that combine them for non-specialists. There is a need for an analytics architecture that (a) takes an uploaded file through validation, cleaning, exploration, dashboards, questions, modelling and reporting in one workflow, (b) computes every number programmatically and uses the language model only to explain, (c) exposes cleaning and clustering decisions so the user can review them, (d) is evaluated with a protocol that covers numerical correctness, model quality, answer relevance and usability, and (e) protects uploaded data.

4. Conclusion

This review examined ten studies spanning visualization recommendation and dashboards [1], [2], [11], natural-language interfaces and analysis agents [4], [10], data quality and automated machine learning [5], [6], [7], customer segmentation [9], [8] and AI in business intelligence [14]. Together they

show that each ingredient of an automated analytics assistant already exists, but that these ingredients are rarely assembled and tested as a whole. Analyst Buddy AI is an attempt at that assembly: one workflow from upload to PDF report, with calculations done by code and the language model used for explanation.

The proposed platform is therefore aimed at lowering the technical barrier to data analytics while keeping a workflow in which a user can inspect the data, visualise important patterns, ask questions, apply machine-learning techniques and receive a summarised insight from a single system. Future work will implement and test each module on real datasets, covering the accuracy of the chart-recommendation logic, the reliability of the natural-language query pipeline, and the quality of the customer-segmentation and forecasting outputs.

The platform is still a prototype with an evaluation plan rather than results. Future work should add a data-quality report the user can approve, show several cluster-validity measures together, compare forecasting models automatically on time-aware validation data, add anomaly detection and role-based access, and build a benchmark of datasets, questions, expected answers and user ratings so that versions of the platform can be compared fairly.

Acknowledgment

The authors are grateful to the Department of Artificial Intelligence and Data Science for its valuable guidance and continuous support during the development of this work.

REFERENCES :

- [1] M. Vartak, S. Huang, T. Siddiqui, S. Madden, and A. Parameswaran, "Towards Visualization Recommendation Systems," 2015.
- [2] S. Sah, R. Mitra, A. Narechania, A. Endert, J. Stasko, and W. Dou, "Generating Analytic Specifications for Data Visualization from Natural Language Queries using Large Language Models," arXiv preprint arXiv:2408.13391, 2024.
- [3] J. Salminen, M. Mustak, M. Sufyan, and B. J. Jansen, "How can algorithms help in segmenting users and customers? A systematic review and research agenda for algorithmic customer segmentation," *Marketing Analytics*, 2024.
- [4] M. Liu, X. Wang, J. Xu, W. Yi, and O. Wolfson, "A systematic review of natural language interfaces for databases," *Frontiers of Computer Science*, 2025.
- [5] M. Baratchi et al., "Automated Machine Learning: Past, Present and Future," *Artificial Intelligence Review*, 2024.
- [6] K. Sha, Y. Zhou, J. Ding, and G. Yang, "Data Quality Dimensions and Tools for Machine Learning," 2024.
- [7] F. Hutter, L. Kotthoff, and J. Vanschoren, "Automated Machine Learning: Methods, Systems, Challenges," Springer, 2025.
- [8] J. Salminen, M. Mustak, and M. Sufyan, "A Systematic Review and Research Agenda for Algorithmic Customer Segmentation," 2023.
- [9] R. Sulthana, S. Ramesh, and A. Jaiswal, "Customer Segmentation Using Machine Learning," 2024.
- [10] W. Xu, Y. Mao, and C. Zhang, "A Relational Database-Driven Data Analysis Report Generation Agent," 2025.
- [11] S. Grant, S. Trevino, and L. Perdue, "Creating Interactive Data Dashboards for Evidence Synthesis," *Cochrane*, 2025.
- [12] S. Sah, R. Mitra, J. Stasko, and W. Dou, "Generating Analytics Specifications for Data Visualization from Natural Language Queries using Large Language Models," 2024.
- [13] M. Liu, X. Wang, J. Xu, W. Yi, and O. Wolfson, "A Systematic Review of Natural Language Interfaces for Databases," *Frontiers of Computer Science*, 2025.