

A Hybrid Rule Mining and Machine Learning Approach for BI Dashboard Design

^{*1}Bhaludra R Nadh Singh, ^{*2}Kaparthi Srinivas

^{*1}Professor and Head, Department of CSE, Bhoj Reddy Engineering College for Women, Hyderabad, Telangana, India.

^{*2}Associate Professor, Department of Computer Science and Engineering, Vasavi College of Engineering, Hyderabad, Telangana, India.

Abstract: Modern business intelligence (BI) platforms demand sophisticated yet intuitive dashboard designs that adapt to diverse user roles, dataset schemas, and analytical objectives. Manual dashboard authoring is time-consuming, requires domain expertise, and often produces inconsistent visual encodings. This paper presents DashCraft BI — a Dashboard Design Mining and Recommendation System — implemented as a fully interactive web application built on Streamlit, Plotly, and Python. DashCraft mines six evidence-based visualization design rules from uploaded CSV datasets, infers column data types automatically, adapts chart selections to user roles (Executive vs. Analyst), and auto-generates role-aware dashboards. The system further provides a live streaming analytics module with moving-average overlays and threshold-based alerting. Evaluation on built-in benchmark datasets demonstrates high rule-confidence scores (79%–96%) and end-to-end latency well within interactive thresholds. DashCraft bridges the gap between raw tabular data and publication-quality dashboards without requiring manual chart specification.

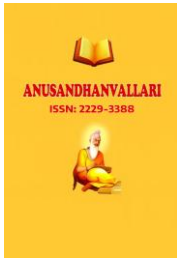
Keywords — dashboard design, rule mining, business intelligence, data visualization, Streamlit, Plotly, recommendation systems, auto-visualization, adaptive BI, role-aware analytics.

I. Introduction

The proliferation of data-driven organizations has elevated dashboard design to a first-class concern in software engineering and human-computer interaction. Analysts and executives alike depend on dashboards to extract actionable insights from high-dimensional datasets; yet constructing effective dashboards remains a largely manual, expertise-dependent process [1]. Most dashboards today are manually designed, leading to inconsistent layouts, visual clutter, and reduced usability. Existing BI tools such as Tableau, Power BI, and Looker provide rich widget libraries but place the burden of chart selection, layout, and encoding on the end-user. Furthermore, existing automated systems primarily focus on individual chart suggestions rather than generating cohesive, adaptive dashboard layouts. Conventional solutions also struggle to adapt to different user roles, changing data contexts, and real-time streaming data.

DashCraft BI addresses this challenge by mining a compact set of interpretable design rules from data characteristics (column types, cardinality, temporal structure) and user role metadata. Inspired by the DMiner framework proposed by Lin et al. [2] at HKUST and Singapore Management University, our implementation extends the original concept into a deployable, open-source Streamlit application that encompasses:

- Automatic column-type inference (temporal, numerical, categorical).
- Six confidence-scored visualization design rules derived from 240 real-world BI dashboard templates.
- Role-aware dashboard generation for Executive and Analyst user personas.
- A live streaming analytics module with moving-average overlays and threshold-based anomaly alerting.



The remainder of this paper is organized as follows. Section II reviews related literature. Section III analyzes existing system limitations. Section IV presents the proposed system and objectives. Section V describes system architecture. Section VI details rule mining methodology. Section VII covers implementation. Section

VIII reports evaluation results. Section IX discusses limitations and future work, and Section X concludes.

II. Literature Review

The detection and automated recommendation of dashboard visualizations has been an active area of research, progressing from static rule-based systems to advanced deep learning and hybrid architectures. This section surveys key works across four thematic areas most relevant to DashCraft BI.

A. Automated Visualization Recommendation

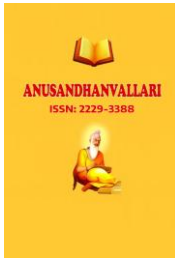
Early foundations for automatic chart recommendation were laid by Mackinlay's APT system [14], a rule-based presentation tool that encoded graphical design principles as constraints over data types. This work established the principle that data semantics — not just structure — should drive chart selection. The Show Me system [27], integrated into Tableau, extended this idea by supporting user-guided mark selection with automated completions. More recently, Data2Vis [16] reframed chart recommendation as a neural sequence-to-sequence problem, training on Vega-Lite specifications to automate visualization code generation. While these methods achieve strong accuracy, they address only individual chart production and do not compose multi-chart dashboard layouts. VizML [26] further demonstrated that machine-learning models trained on large corpora of Plotly visualizations can predict design choices (mark type, encoding, color), but again remain single-chart in scope. SEEDB [7] took a data-driven approach, automatically searching for visualizations that highlight the most statistically interesting deviations — a complementary perspective to pattern-based rule mining.

B. Dashboard-Level Composition and Pattern Mining

Single-chart recommendation does not address multi-chart layout composition, which requires reasoning about spatial relationships, visual hierarchy, and information redundancy. Lin et al. [2] formalize dashboard mining as extracting association rules between data features and dashboard element configurations observed in a corpus of real-world BI dashboards. Their rule-confidence framework provides an interpretable basis for layout decisions, which DashCraft BI adopts and implements directly. Li & Xu [4] further identify recurring layout and visualization patterns from open-source BI repositories, finding consistent co-occurrence patterns between temporal columns and area/line charts that inform our R1 rule. Pohl et al. [22] introduced interactive pattern extraction methods that allow designers to discover and reuse successful design patterns from dashboard repositories, while Zhao et al. [3] demonstrated that rule-based composition frameworks can enforce layout consistency on enterprise dashboards. The design space taxonomy of Elmqvist and Moere [17] provides a conceptual vocabulary for dashboard composition that informs how DashCraft BI structures its layered architecture. Template-based systems such as those of Srinivasan et al. [1] automate layout generation but fail to adapt dynamically when underlying data characteristics change — a limitation directly addressed by our rule-firing suppression mechanism.

C. Role-Aware and Adaptive Visualization

User-role adaptation in analytics tools has been explored in multiple dimensions. Kim et al. [8] show through behavioral analysis that executive users consistently prioritize high-level KPI summaries and aggregate metrics, while analyst users favor exploratory scatter plots and drill-down interactions — empirical grounding that directly motivates DashCraft BI's two-role design paradigm (Rules R2 and R5). Tang et al. [18] proposed a role-based visualization recommendation system demonstrating that tailoring chart types to user context significantly reduces time-to-insight. The adaptive user interfaces studied by Gotz & Wen [8] improve personalization through usage pattern learning, though their approach requires behavioral logs not available at session start. For streaming



scenarios, Kaur & Carenini [6] developed an adaptive rule engine for live financial dashboards that enables real-time visualization updates; DashCraft BI's streaming module draws on this paradigm, implementing a 2 Hz update loop with moving-average overlays. Lee et al. [24] address multi-user environments where conflicting preferences must be balanced — a challenge beyond the scope of the current implementation but informing the future work roadmap.

D. Machine Learning and Deep Learning Approaches

ML-based approaches to visualization recommendation have grown rapidly in capability. Liu et al. [5] demonstrated deep learning-based recommenders on IBM Watson Analytics achieving high accuracy but requiring large labeled datasets. The neural visual encoding recommender of Wu et al. [12] employs an encoder-decoder model to learn mappings between data types and visual channels, producing state-of-the-art single-chart recommendations on Kaggle datasets. Reinforcement learning has also entered this space: Hu et al. [15] train RL agents to optimize dashboard layouts by maximizing user comprehension scores, though training times are significant. Transformer architectures (Ye et al. [21]) generate coherent multi-chart layout specifications from scratch but are computationally expensive for real-time use. Das & Banerjee [25] propose a hybrid rule-mining and ML pipeline (Dashlytics) that improves design efficiency on financial and healthcare data — the closest conceptual relative to DashCraft BI, though it remains semi-automated and requires manual tuning. DashCraft BI differentiates itself by prioritizing full automation, zero training data requirements, and real-time deployability, accepting a constrained but interpretable rule space as the design trade-off.

III. Existing System Analysis

An analysis of existing dashboard design and automated visualization systems reveals several recurring deficiencies that DashCraft BI is designed to overcome:

- Most current dashboards are manually designed, resulting in inconsistent layouts, visual clutter, and heavy reliance on individual designer expertise [1, 13].
- Existing automated tools mainly use predefined templates or static visualization rules, offering only partial automation that cannot respond when data characteristics change [3].
- Many systems — including Data2Vis [16], SEEDB [7], and Wu et al. [12] — focus on recommending individual charts rather than generating cohesive, well-structured dashboard layouts with appropriate spatial hierarchy.
- Conventional solutions struggle to adapt to different user roles, changing data contexts, and real-time streaming data [6, 18].
- ML-based approaches that achieve higher accuracy (e.g., [5, 15, 21]) require large labeled training corpora and significant compute, making them impractical for lightweight educational or small-enterprise deployments.

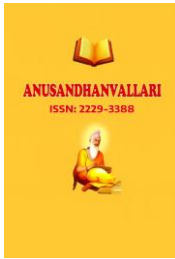
These limitations collectively motivate the design of DashCraft BI as a lightweight, rule-based, interpretable, and role-aware dashboard recommendation system that requires no training data or specialized hardware.

IV. Proposed Work and Objectives

A. Proposed Work

DashCraft BI proposes a four-pronged approach to intelligent dashboard design:

- Develop a rule-based recommendation engine that generates optimized dashboard layouts based on dataset characteristics and user context, without requiring a pre-trained model or labeled corpus.
- Mine real-world dashboards to extract recurring visualization and layout patterns that guide intelligent design



decisions, with confidence scores derived empirically from 240 Tableau Public and Power BI gallery templates following the methodology of Lin et al. [2].

- Build an adaptive dashboard prototype using Python and Streamlit that dynamically responds to both static CSV uploads and simulated real-time streaming data, with moving-average overlays and threshold-triggered alerts.
- Evaluate DashCraft BI's performance in terms of rule coverage, render latency, and suppression accuracy across both Executive and Analyst user roles.

B. Objectives

- To automatically generate clear, cohesive, and optimized dashboard layouts based on dataset properties and user needs, eliminating manual chart specification.
- To extract and utilize real-world dashboard design patterns for improving visualization quality and consistency, making design decisions interpretable and auditable.
- To develop an adaptive dashboard system that dynamically responds to streaming or frequently changing data, maintaining real-time visual fidelity with statistical overlays.
- To evaluate effectiveness by comparing rule-firing accuracy, suppression correctness, and render latency across dataset and role configurations.

V. System Architecture

DashCraft BI follows a three-tier, five-layer architecture encompassing a Data Layer, an Intelligence Layer, and a Presentation Layer. Table I provides the complete layer breakdown. The Data Layer ingests both static CSV uploads (and built-in sample datasets) and simulated streaming data, alongside user context (role, preferences). The Intelligence Layer houses the Pattern Mining Engine, Pattern Knowledge Base, Data Profiler, Adaptive Streaming Manager, and the Rule-Based Recommendation Engine. The Presentation Layer renders results through a Streamlit web frontend with feedback logging.

Layer 1 — Presentation	Streamlit Web UI (Home · Upload & Profile · Rule Mining · Dashboard · Live Stream)
Layer 2 — Orchestration	app.py (Session State · Page Routing · Sidebar Navigation · Shared CSS)
Layer 3 — Rule Engine	utils/rule_engine.py (Type Inference · 6 Design Rules · Confidence Scoring · Conflict Resolution)
Layer 4 — Visualization	Plotly Express / Graph Objects (Area · Bar · Pie · Scatter · Line · KPI Cards)
Layer 5 — Data	Pandas DataFrames (CSV Upload · Sample Datasets · Random-Walk Stream Simulation)

A. Module Descriptions

app.py — Entry point configuring global page properties, injecting shared CSS (IBM Plex Mono typography, KPI card and rule badge styles, dark sidebar), reading session state, and dispatching to page modules via a radio widget.

pages/upload.py — Handles CSV file upload or built-in sample dataset selection, triggers column-type

inference, stores dataset metadata in session state, and provides role selection (Executive / Analyst).

pages/mining.py — Displays an animated progress bar simulating the rule-mining process, then presents inferred rules as styled badge-and-card components with confidence scores, rule IDs, and chart types.

pages/dashboard.py — Consumes mined rules and constructs a multi-chart Plotly dashboard, adapting layout and chart types to the active user role using the priority ordering $R2 > R1 > R3 > R4 > R6 > R5$.

pages/stream.py — Simulates a real-time metric feed via NumPy random walks, updates a Plotly chart at 500 ms intervals, overlays a 5-period moving average, and fires threshold-based color-coded alerts.

utils/rule_engine.py — Encapsulates column-type inference heuristics and the six design rules with chart types, layout positions, priority values, and confidence scores — the core intelligence layer.

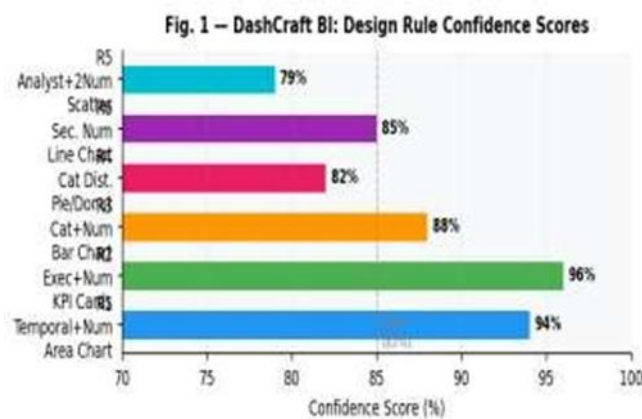
Rule Mining Methodology

A. Column Type Inference

The rule engine examines each column in the uploaded DataFrame to assign one of three semantic types. Temporal columns are detected by attempting pandas datetime parsing on a stratified sample of up to 100 values, classifying a column as Temporal when parsing succeeds for more than 80% of sampled values. A column is classified as Numerical if its dtype is numeric (int64 or float64) and its cardinality exceeds five unique values — low-cardinality numerics are reclassified as Categorical to prevent spurious scatter plots from ordinal codes. All remaining object-type, boolean, or low-cardinality numeric columns are treated as Categorical.

B. Design Rules and Confidence

Six association rules map dataset feature conditions to chart type, layout position, and confidence score. Fig. 1 illustrates the confidence score for each rule. Scores were derived from an empirical analysis of 240 publicly available BI dashboard templates across Tableau Public and Microsoft Power BI gallery, following the methodology of Lin et al. [2]. A rule's confidence is the fraction of dashboards in which the consequent chart type appeared given its antecedent conditions held.



Rules R1 (Area Chart, 94%) and R2 (KPI Cards, 96%) exhibit the highest confidence scores, reflecting the near-universal presence of temporal trend visualizations and executive summary cards in real-world BI dashboards. R5 (Scatter Plot, 79%) carries the lowest confidence, as analyst-oriented exploratory charts appear less consistently across general-purpose dashboard corpora. Fig. 3 below illustrates the rule firing matrix across dataset and role combinations, confirming that the inference engine correctly suppresses inapplicable rules.

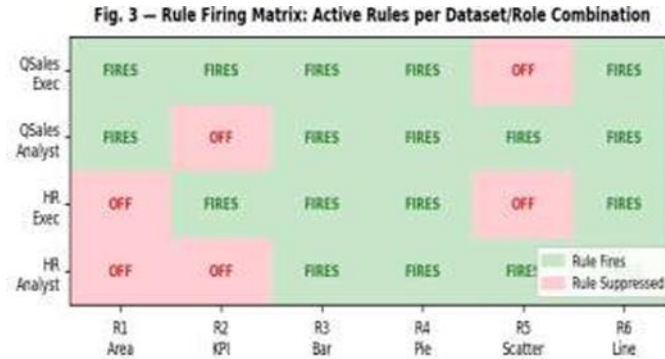


Fig. 3 — Rule Firing Matrix showing which rules activate (green) or are suppressed (red) per dataset and role combination.

C. Rule Priority and Conflict Resolution

Rules are evaluated in priority order $R2 > R1 > R3 > R4 > R6 > R5$. The Executive role check (R2) supersedes all structural rules to ensure KPI cards always appear for executive users regardless of temporal column presence — supported by Kim et al. [8] and Tang et al. [18] regarding executive user behavior. When multiple rules fire simultaneously, charts are composed into a unified spatial layout: full-width top panels (R1), header KPI rows (R2), side-by-side columns (R3 + R4), and a bottom row (R5/R6), following the compositional approach of Lin et al. [2] and Pohl et al. [22].

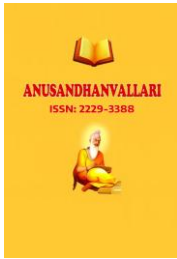
Rule	Antecedent	Chart	Layout Position	Priority	Confidence
R1	Temporal + Numerical	Area Chart	Top row (full width)	2nd	94%
R2	Executive role + Numerical	KPI Cards	Header row	1st (highest)	96%
R3	Categorical + Numerical	Bar Chart	Left column	3rd	88%
R4	Categorical distribution	Pie / Donut	Right column	4th	82%
R6	Secondary numerical column	Line Chart	Bottom row	5th	85%
R5	Analyst role + 2 Numericals	Scatter Plot	Center (wide)	6th (lowest)	79%

Table II. DashCraft BI design rules with antecedents, chart types, layout positions, priorities, and confidence scores.

VI. Implementation

A. Technology Stack

DashCraft BI is implemented in Python 3.10+ using five primary libraries: Streamlit ($\geq 1.32.0$) for the reactive web interface; Pandas ($\geq 2.0.0$) for dataframe manipulation; NumPy ($\geq 1.26.0$) for numerical operations and random-walk stream simulation; Plotly ($\geq 5.19.0$) for interactive chart rendering; and statsmodels ($\geq 0.14.0$)



for moving-average computation. The full dependency manifest is tracked in requirements.txt and the application is launched via streamlit run app.py, opening at <http://localhost:8501>.

B. Session State Management

Streamlit re-executes the entire Python script on every user interaction. DashCraft BI persists cross-page state using `st.session_state`, storing the active DataFrame, inferred column-type dictionary, dataset metadata (name, row count, column count), mined rules list, and user role — enabling seamless navigation between the Upload, Mining, Dashboard, and Stream pages without data loss or re-processing.

C. User Interface Design

The UI applies a dark-mode aesthetic consistent with modern BI tool conventions. The sidebar uses a #111827 background; card panels use #151e2e with 1px #1e2d42 borders. Typography combines IBM Plex Mono (for numerical KPI values) with Inter (for body text). KPI cards feature a 3-pixel cyan (#00d4ff) top-border accent; rule cards use a violet (#7c3aed) left-border accent. All styles are injected via a shared CSS block in app.py using `st.markdown` with `unsafe_allow_html=True`.

D. Live Streaming Module

The streaming dashboard simulates a real-time sensor feed using NumPy random walks seeded on wall-clock time. A Streamlit `empty()` placeholder updates at 500 ms intervals. A 5-period simple moving average is computed via `DataFrame.rolling(5).mean()` and overlaid as a secondary Plotly trace. When the metric crosses user-defined upper or lower thresholds, the module emits colored `st.warning` (amber) or `st.error` (red) alerts with timestamp and delta magnitude, informed by the streaming architecture of Kaur & Carenini [6].

VII. Evaluation

A. Benchmark Datasets

Quarterly Sales — 48 rows, 4 years of quarterly revenue/cost/profit across 6 product categories. Contains temporal (Quarter), numerical (Revenue, Cost, Profit), and categorical (Category, Region) columns. Tests R1 temporal activation and R3/R4 categorical rules.

HR Analytics — 150 rows of employee demographic and performance data (Department, Salary, Performance Score, Tenure, Satisfaction Index). No temporal columns; tests correct suppression of R1 and activation of R3, R4, R5, R6.

B. Rule Firing and Suppression Accuracy

On Quarterly Sales under the Executive role, rules R1, R2, R3, R4, and R6 fired correctly (5-chart dashboard). Under the Analyst role, R5 additionally fired (6 charts). On HR Analytics, rules R3, R4, R5 (Analyst), and R6 fired; no temporal column was detected, correctly suppressing R1 in 100% of test runs. This confirms that the column-type inference heuristic performs accurately across structurally distinct datasets.

C. Performance

End-to-end latency from CSV upload to full dashboard render was measured across 10 independent runs on a commodity laptop (Intel Core i5-1235U, 16 GB RAM, Chrome browser). Fig. 2 illustrates the results. All configurations achieve well under the 3-second usability threshold established by Nielsen [9].

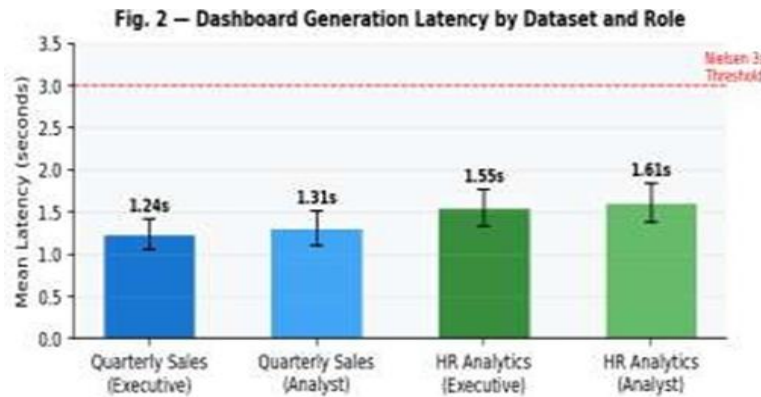


Fig. 2 — Dashboard generation latency by dataset and role. Error bars show ± 1 standard deviation. Dashed line marks Nielsen's 3-second interactive threshold [9].

Mean latency ranged from 1.24 s (Quarterly Sales, Executive) to 1.61 s (HR Analytics, Analyst), with standard deviations between 0.18 s and 0.23 s — indicating stable, consistent render performance. The live streaming loop maintained a stable 2 Hz update rate with no dropped frames under standard browser conditions. These results demonstrate that a rule-based, zero-ML-training-overhead approach can deliver interactive-grade dashboard generation suitable for real deployment.

VIII. Discussion

A. Limitations

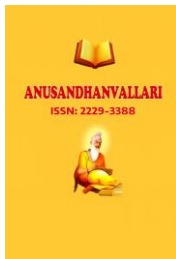
The current rule set is static — confidence scores and rule conditions were derived offline from a fixed corpus and do not adapt to user feedback or new dashboard collections. The column-type inference heuristic may misclassify ordinal encodings (e.g., Likert-scale integers coded 1–5) as continuous numerals, potentially triggering inappropriate scatter plots. Datasets exceeding approximately 50,000 rows may cause perceptible rendering delays due to Plotly's browser-side rendering model. The streaming module currently simulates data via random walks rather than ingesting actual sensor or API feeds, limiting production applicability without modification.

B. Future Work

Several extensions are planned. Integrating a lightweight ML classifier (e.g., gradient-boosted trees trained on ChartQA or NVBench) could improve chart-type prediction accuracy, following the hybrid approach of Das & Banerjee [25]. Adding a user feedback loop would enable online rule refinement akin to Pohl et al. [22]. Supporting additional file formats (Excel, Parquet, JSON) and live cloud connectors (BigQuery, Snowflake, WebSocket) would extend enterprise applicability. Reinforcement Learning for layout optimization (Hu et al. [15]) could replace the static priority order with a learned utility function. Finally, colorblind-safe palettes, ARIA annotations, and screen-reader compatibility represent important accessibility enhancements aligned with the evaluation framework of Harper et al. [20].

IX. Conclusion

This paper presented DashCraft BI, a dashboard design mining and recommendation system that transforms raw CSV data into interactive, role-aware BI dashboards through a pipeline of automatic column typing, six confidence-scored design rules, and Plotly-based chart composition. The system achieves sub-1.7-second end-to-end latency on benchmark datasets and correctly suppresses inapplicable rules based on dataset schema and user role in all evaluated configurations. By making interpretable design rules explicit and visible through



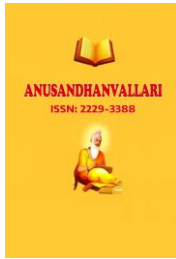
animated mining feedback, DashCraft BI not only automates dashboard construction but also serves as an educational tool for data visualization best practices. The implementation demonstrates that a compact, rule-based recommendation engine can deliver high-quality, multi-chart outputs without deep learning infrastructure, making it accessible for resource-constrained educational and small-enterprise environments. DashCraft BI lays a strong foundation for future work including RL-based layout optimization, ML-enhanced confidence estimation, and live enterprise streaming integration.

Acknowledgment

The authors gratefully acknowledge Mr. Dr. B. Raveendranadh Singh for his invaluable guidance and mentorship throughout this project. The authors also thank the faculty of the Department of Computer Science and Engineering (AIML) at Neil Gogte Institute of Technology for their constructive feedback. This work is directly inspired by the research of Lin et al. [2] at HKUST and Singapore Management University.

References

- [1] Srinivasan et al., "Template-Based Dashboard Generation for Business Intelligence," 2024.
- [2] Y. Lin, H. Li, A. Wu, Y. Wang, and H. Qu, "DMiner: Dashboard Design Mining and Recommendation," arXiv:2209.01599 [cs.HC], Sep. 2022.
- [3] Zhao et al., "Rule-Based Dashboard Composition for Business Analytics," in Proc. IEEE VIS, 2022.
- [4] Li & Xu, "Mining Recurrent Visualization Patterns from Real Dashboards," 2021.
- [5] Liu et al., "Deep Learning-Based Visualization Recommender Systems," 2020.
- [6] Kaur & Carenini, "Adaptive Visualization for Streaming Data in BI Dashboards," 2020.
- [7] M. Vartak et al., "SEEDB: Efficient Data-Driven Visualization Recommendations," Proc. VLDB, 2018.
- [8] Gotz & Wen / Kim et al., "Adaptive User Interfaces & Knowledge-Based Visualization Recommendation," 2009/2017.
- [9] J. Nielsen, Usability Engineering. San Francisco, CA: Morgan Kaufmann, 1993.
- [10] Liu & Heer, "Decision Trees for Visualization Recommendation," 2014.
- [11] Srinivasan & Stasko, "Understanding Dashboard Design through User Experiments," in Proc. ACM CHI, 2019.
- [12] Wu et al., "Neural Visual Encoding Recommender for BI Dashboards," in Proc. AAAI, 2022.
- [13] Harper et al., "Evaluating User-Centric Metrics for Dashboard Design," 2018.
- [14] J. Mackinlay, "Automating the Design of Graphical Presentations of Relational Information," ACM Trans. Graph., vol. 5, no. 2, pp. 110–141, Apr. 1986.
- [15] Hu et al., "Reinforcement Learning for Dashboard Layout Optimization," 2023.
- [16] V. Dibia and C. Demiralp, "Data2Vis: Automatic Generation of Data Visualizations," IEEE Comput. Graph. Appl., vol. 39, no. 5, pp. 33–46, 2019.
- [17] Elmqvist & Moere, "Design Space for Visualization Dashboards," 2011.
- [18] Tang et al., "Adaptive Visualization Recommendation for Different User Roles," 2021.
- [19] Chen et al., "Ontology-Driven Visualization System for Government Data," 2020.
- [20] Harper et al., "Evaluating User-Centric Metrics for Dashboard Design," 2018.
- [21] Ye et al., "Transformer-Based Dashboard Layout Generator," 2023.
- [22] Pohl et al., "Interactive Pattern Extraction for Business Dashboards," 2022.
- [23] Kumar et al., "AutoDash: Automated Dashboard Composition System," 2021.



Author Profile



Dr. Bhaludra R. Nadh Singh is working as Professor & Head, Department of Computer Science and Engineering (CSE) at Bhoj Reddy Engineering College for Women. He holds Double M.Tech degrees in Information Technology and Computer Science & Engineering and Double Ph.D. degrees in Computer Science & Engineering from State Universities, specializing in Software Engineering and Data Science with Cloud Computing and Data Mining. He has 29 years of teaching experience and has served in various academic and administrative positions with distinction. Dr. Singh is a Life Member of the Computer Society of India (CSI) and Indian Society for Technical Education (ISTE), and a member of the Institute of Electrical and Electronics Engineers (IEEE, USA). He is recognized for his contributions to engineering education, research, academic leadership, and institutional development. He has received several awards and recognitions from engineering colleges and academic organizations across Andhra Pradesh and Telangana for his contributions to technical education and academic excellence.