

From Chaos to Clarity: Data Visualisation Approaches for High-Demand Data

Ram Niwas¹; Yogesh Deshmukh¹; Sneha Vadavi²; Khushi Chaudhary²

^{1,2}Department of Master of Computer Applications
Alard School of Business Management, Pune, India

Publication Date: 2026/04/01

Abstract: The exponential growth of digital data has created what may be termed informational chaos — a condition in which the volume, velocity, and variety of incoming data exceed the processing and interpretive capacity of conventional analytics systems. Organisations operating in high-demand data environments, particularly in domains such as e-commerce, finance, and healthcare, face mounting pressure to derive real-time, actionable insights from massive and continuously evolving datasets. Despite the proliferation of visualisation tools, the existing literature offers limited structured guidance on how to select and deploy these tools specifically within high-demand contexts, defined in this paper as environments characterised simultaneously by high data volume (millions or billions of records), high velocity (near real-time ingestion), high variety (heterogeneous data modalities), and elevated decision pressure. This study addresses that gap through three primary contributions. First, a task-oriented taxonomy is proposed, classifying high-demand visualisation activities into four operational categories: Monitoring, Alerting, Exploration, and Explanation. Second, a six-criterion scoring framework is introduced, evaluating tools across Scalability (C1), Latency (C2), Cognitive Load (C3), Literacy Support (C4), Domain-Adaptivity (C5), and Big-Data Integration (C6). Third, the framework is empirically grounded through a primary survey of thirty working IT and analytics professionals, supplemented by an illustrative case study comparing four leading tools — Tableau, Power BI, D3.js, and Plotly — across the e-commerce domain. Survey findings indicate that Power BI is the most widely adopted tool (50% most-frequent usage), with an overall mean effectiveness rating of 8.37 out of 10, while Tableau is preferred for exploratory and narrative visualisation tasks. The framework proposed herein offers a reproducible methodology for practitioners and researchers seeking tool-selection guidance in data-intensive operational settings.

Keywords: Data Visualisation, High-Demand Data, Tool Selection Framework, Dashboard Design, Big Data Analytics, Visualisation Literacy, Cognitive Load, E-Commerce Analytics.

How to Cite: Ram Niwas; Yogesh Deshmukh; Sneha Vadavi; Khushi Chaudhary (2026) From Chaos to Clarity: Data Visualisation Approaches for High-Demand Data. *International Journal of Innovative Science and Research Technology*, 11(3), 2859-2867. <https://doi.org/10.38124/ijisrt/26mar1502>

I. INTRODUCTION

Contemporary organisations generate data at an unprecedented and still-accelerating rate. By current estimates, the global data sphere is expected to surpass 175 zettabytes by 2025, with a significant portion of this data produced in environments requiring near-instantaneous analytical interpretation. Within this context, data visualisation has emerged not merely as a convenience but as a strategic necessity — a discipline that transforms raw, high-dimensional data streams into perceptually accessible representations capable of informing rapid decisions.

Gorodov and Gubarev [1] characterise the foundational challenge of big data visualisation through the lens of the 3V model — Volume, Velocity, and Variety — arguing that conventional tabular and static-chart representations are structurally inadequate for data exhibiting these three properties simultaneously. Their analysis situates the core problem not as one of computational capacity alone, but as a perceptual and cognitive problem: when the rate and scale of data generation outpaces human cognitive processing, analytical tools must mediate complexity rather than merely display it. This argument provides the theoretical basis for the concept of informational chaos, as employed in the present paper — a state in which data abundance paradoxically impedes rather than enables effective decision-making.

Complementing this perspective, Börner, Bueckle, and Ginda [2] document a growing global awareness of the need for data visualisation literacy (DVL) — the capacity of individuals not only to consume visual data representations but to critically evaluate and construct them. Their longitudinal study across multiple educational and professional contexts reveals that DVL competencies are unevenly distributed and often underdeveloped even among data professionals, a finding with direct implications for tool design and selection in high-demand environments. If the intended audience of a visualisation dashboard lacks the interpretive competencies to engage with its outputs, the sophistication of the underlying rendering engine is inconsequential.

The intersection of these two challenges — the technical demands of high-demand data and the human-centred demands of visualisation literacy — defines what this paper identifies as the central problem in enterprise analytics: the absence of a structured, multi-criteria framework for selecting data visualisation tools that addresses both dimensions simultaneously. Existing tool comparison studies, such as the comparative analysis by Geethu C. Nair [8] and the scholarly visualisation survey by Liu et al. [5], tend to evaluate tools in isolation or along single dimensions such as feature richness or rendering performance. They do not provide a systematic methodology calibrated to the specific operational demands of high-velocity, high-volume, heterogeneous data environments under decision pressure.

This paper addresses that gap through three connected contributions: (1) a task-oriented taxonomy of high-demand visualisation activities; (2) a six-criteria scoring framework for tool selection; and (3) an empirical validation of this framework through primary survey data collected from thirty working IT and analytics professionals, supplemented by an illustrative case study in e-commerce customer behaviour analytics. The remainder of this paper is structured as follows. Section 2 reviews the relevant background literature. Section 3 states the research problem and formulates four research questions. Section 4 summarises the paper's contributions. Section 5 describes the methodology. Sections 6 and 7 detail the taxonomy and evaluation criteria respectively. Section 8 presents the case study and survey findings. Section 9 discusses findings and limitations. Section 10 concludes with directions for future work.

II. BACKGROUND AND RELATED WORK

➤ *Data Visualisation Literacy*

Börner, Bueckle, and Ginda [2] present the Data Visualisation Literacy Framework (DVL-FW), a structured competency model derived from empirical studies across K-12 education, higher education, and professional contexts. The framework decomposes visualisation literacy into four sequential competencies: identifying the chart type, interpreting the data encoding, executing appropriate analytical operations, and inferring meaning from the resulting representation.

Crucially, their study reveals that many practitioners, even in data-rich organisations, are proficient only at the first two stages, and struggle with the inferential demands of the latter two. The paper leaves open the question of how tool design can compensate for these literacy gaps in operational settings, particularly those involving high data throughput and time-constrained decision-making. The present study incorporates DVL considerations directly into the Literacy Support criterion (C4) of the proposed scoring framework, treating it as a measurable tool attribute rather than a user characteristic.

➤ *Dashboard Meta-Models and User-Centred Design*

Vazquez-Ingelmo et al. [3] propose a meta-model for dashboard design, developed through a series of focus groups with both business and technical stakeholders. Their work demonstrates that dashboard effectiveness is fundamentally contextual: the same information architecture that serves a data-literate analyst navigating exploratory queries may actively impede a business executive seeking a single, actionable performance indicator. The study identifies a persistent misalignment between the design assumptions embedded in general-purpose dashboarding tools and the actual cognitive workflows of diverse user populations. However, the meta-model is primarily prescriptive at the design level and does not offer a comparative tool-selection mechanism. This paper extends that line of inquiry by translating user-centred design principles into quantifiable evaluation criteria, enabling systematic tool comparison rather than post-hoc design critique.

➤ *Surveys of Visualisation Tools and Techniques*

Harsha Vardhan and Sree Devi [4] conduct a broad survey of visualisation techniques applicable to large-scale datasets, categorising existing approaches by data type, interaction modality, and rendering mechanism. Their analysis is comprehensive but operates at the level of technique rather than tool, and does not account for the operational constraints — latency requirements, concurrent user load, or integration with upstream data pipelines — that distinguish production analytics environments from research prototypes. Similarly, Liu et al. [5] survey scholarly data visualisation from an academic research perspective, identifying representational innovation as the primary evaluative criterion. Neither study foregrounds the practitioner's tool-selection decision. The current paper addresses this practitioner-oriented gap by combining a review of tool capabilities with empirical feedback from professionals actively using these tools in large-dataset environments.

Geethu C. Nair [8] offers a more recent comparative analysis of contemporary visualisation tools, examining Tableau, Power BI, D3.js, and others along dimensions of usability, performance, and learning curve. While this work is directly relevant to the current paper's case study, its evaluation methodology is largely qualitative and does not employ a multi-criteria scoring model. The present study formalises and operationalises the comparison through the six-criteria framework described in Section 7, and supplements it with primary survey evidence.

➤ *Big Data Visualisation and Domain Applications*

Gorodov and Gubarev [1] provide an analytical review of visualisation methods for big data, situating the problem within information overload theory and arguing for the necessity of interactive, layered, and algorithmically adaptive visual representations. Their review establishes the technical vocabulary — particularly the 3V characterisation — that underpins the concept of high-demand data as employed in this paper. Khade [16] extends this discussion into an applied context, demonstrating the use of D3.js alongside big data pipelines for e-commerce customer behaviour analytics. That study illustrates both the expressive power of programmatic visualisation and the considerable technical overhead associated with its deployment in production settings. Dwyer et al. [6] further argue, from the perspective of immersive analytics, that the challenges of high-demand visualisation are not solvable through incremental improvement of existing two-dimensional dashboard paradigms, but may require fundamentally different perceptual and interactional modalities. Collectively, these works establish the domain problem that the present paper seeks to address with a structured framework.

III. PROBLEM STATEMENT AND RESEARCH QUESTIONS

Despite significant advances in both visualisation technology and the theoretical understanding of visualisation literacy, the field lacks a unified, criterion-based framework for selecting visualisation tools specifically in high-demand environments. Practitioners in IT and analytics roles — particularly those operating at the intersection of large-scale data pipelines and business intelligence — routinely face tool-selection decisions without the benefit of structured guidance calibrated to operational constraints such as query latency, concurrent user load, and the cognitive accessibility requirements of mixed-literacy audiences. This gap is consequential not only for the quality of analytical outputs but for organisational decision-making capacity, particularly in time-sensitive domains such as e-commerce, financial trading, and clinical operations management.

For MCA graduates and IT professionals entering the workforce, this gap is especially acute: existing academic curricula often introduce visualisation tools through feature demonstrations rather than through needs-based selection methodologies, leaving professionals ill-equipped to justify tool choices in procurement or architecture review contexts. The following four research questions guide the present investigation:

- RQ1: What are the primary categories of visualisation tasks that characterise high-demand data environments?
- RQ2: What criteria should be employed to evaluate and compare data visualisation tools in high-demand contexts?
- RQ3: How do leading visualisation tools — specifically Tableau, Power BI, D3.js, and Plotly — perform against these criteria in an e-commerce analytics context?

- RQ4: How do the perceptions of working IT and analytics professionals validate or challenge the framework-derived tool evaluations?

IV. CONTRIBUTIONS

This paper makes three principal contributions to the data visualisation literature:

➤ *Contribution 1 — Task-Oriented Taxonomy:*

A four-category taxonomy of high-demand visualisation tasks (Monitoring, Alerting, Exploration, Explanation) is proposed and elaborated with real-world examples. Unlike existing tool surveys, which classify tools by feature set or rendering technology, this taxonomy classifies by operational intent, enabling a need-first approach to tool selection that is absent from the current literature.

➤ *Contribution 2 — Six-Criteria Scoring Framework:*

A replicable scoring model employing six criteria (C1–C6) is developed, enabling systematic comparison of visualisation tools across dimensions that are simultaneously technical (scalability, latency, big-data integration) and human-centred (cognitive load, literacy support, domain adaptivity). Existing comparative studies evaluate these dimensions qualitatively or in isolation; this framework integrates them into a single, quantifiable evaluation instrument.

➤ *Contribution 3 — Empirical Validation:*

The framework is validated through primary survey data collected from thirty working IT and analytics professionals spanning multiple roles and experience levels. This empirical grounding distinguishes the present study from purely theoretical frameworks or illustrative case studies, providing practitioner-anchored evidence for the relative positioning of tools within the scoring model.

V. METHODOLOGY

➤ *Research Design*

This study employs a mixed-methods research design, combining a systematic literature review, framework development, illustrative case study analysis, and primary survey-based empirical validation. The systematic review serves as the evidentiary foundation for the taxonomy and scoring criteria. The case study grounds the framework in a concrete domain scenario. The survey provides practitioner-level validation of framework-derived tool scores.

➤ *Systematic Literature Review Protocol*

The literature review was conducted following a PRISMA-inspired protocol. In the identification phase, candidate studies were retrieved from Google Scholar, IEEE Xplore, and Proceidia/Elsevier databases using search terms including 'data visualisation tools comparison', 'big data visualisation survey', 'dashboard design framework', 'visualisation literacy', and 'high-demand analytics'. The search

was not date-restricted, given the relevance of foundational works to the conceptual framework, though priority was given to studies published after 2013. In the screening phase, abstracts were evaluated for relevance to at least one of the following themes: tool evaluation, visualisation literacy, dashboard design, big data visualisation, or domain-specific analytics applications. Studies presenting empirical findings, structured taxonomies, or comparative analyses were preferred over opinion pieces. In the eligibility phase, eight studies were

selected for full review and citation, as presented in Table 1. Exclusion criteria included papers focusing exclusively on visualisation algorithms without tool-level discussion, papers addressing visualisation solely in non-analytical domains (e.g., scientific simulation), and papers for which full text was inaccessible. The synthesis phase involved thematic analysis of the selected papers, with particular attention to gaps in tool-selection guidance for high-demand environments.

Table 1. Literature Review Summary

Ref.	Authors (Year)	Focus Area	Key Contribution	Gap Addressed in Present Study
[1]	Gorodov & Gubarev (2013)	Big Data vis. methods	3V-based taxonomy of vis. challenges	No tool-selection criteria provided
[2]	Borner et al. (2019)	Visualisation literacy (DVL-FW)	Competency framework across contexts	No integration into tool scoring
[3]	Vazquez-Ingelmo et al. (2024)	Dashboard meta-model	User-centred design via focus groups	Design-level; not a selection framework
[4]	Harsha Vardhan & Sree Devi (2020)	Large dataset vis. techniques	Survey of rendering methods	Technique-level; no tool comparison
[5]	Liu et al. (2018)	Scholarly data visualisation	Review of academic vis. approaches	Research focus; no practitioner criteria
[6]	Dwyer et al. (2020)	Immersive analytics lab	Novel modalities for complex data	Experimental; not for standard BI tools
[7]	Khade (2016)	E-commerce customer analytics	D3.js for big data application	Single-tool; no framework for comparison
[8]	Geethu C. Nair (2024)	Tool comparative analysis	Qualitative comparison of major tools	No multi-criteria scoring model

➤ Framework Development

The taxonomy and scoring framework were developed through a two-stage process. In the first stage, the themes identified in the literature review were abstracted into candidate task categories and evaluation criteria. In the second stage, these candidates were refined and operationalised through review against the survey data, ensuring that the criteria reflected dimensions that practitioners actively consider in their tool-use experiences. Each criterion was assigned a weight and a scoring anchor (1–5 scale) based on the frequency and salience with which related concerns appeared in the professional literature and survey responses.

➤ Primary Survey

A structured questionnaire comprising ten questions was administered to a purposive sample of thirty working IT and analytics professionals. Respondents were selected on the basis of active engagement with large-scale data visualisation in professional contexts, spanning roles including Data Analyst, BI Developer, Data Scientist, Analytics Manager, and Data

Engineer, among others. The survey collected data on tools employed, most-frequent tool choice, perceived strengths and difficulties, self-assessed effectiveness ratings on a 1–10 scale, and tool recommendations for high-demand scenarios. The resulting data were analysed both quantitatively (frequency distributions, mean effectiveness scores) and qualitatively (thematic coding of open-ended responses), and are used in Section 8 to validate the framework-derived tool scores.

VI. TAXONOMY OF HIGH-DEMAND VISUALISATION TASKS

High-demand data environments are not monolithic: the visualisation task faced by a network operations centre engineer monitoring server uptime differs fundamentally from that faced by a retail analyst exploring seasonal purchase patterns. This section introduces a four-category taxonomy that distinguishes high-demand visualisation tasks by their operational intent, time horizon, and analytical depth. The taxonomy is designed to be the first step in the tool-selection process: by identifying

the dominant task type in a given use case, practitioners can meaningfully constrain the space of appropriate tools before applying the scoring framework.

Table 2. Taxonomy of High-Demand Visualisation Tasks

Task Class	Definition	Time Horizon	Representative Industry Example	Key Visualisation Requirements
Monitoring	Continuous tracking of system or business KPIs against predefined thresholds	Real-time to near-real-time	E-commerce: live order fulfilment rate dashboard during peak shopping events such as Black Friday	Low latency, minimal visual complexity, persistent display, auto-refresh
Alerting	Detection and surfacing of anomalies, threshold breaches, or statistically significant deviations	Milliseconds to minutes	Finance: automated flag triggered when transaction volume deviates by more than 3 standard deviations from the hourly mean	Event-driven rendering, colour-coded severity, minimal cognitive overhead for rapid triage
Exploration	Open-ended, hypothesis-driven interaction with high-dimensional data to discover patterns or relationships	Minutes to hours	Healthcare: clinical data scientist examining comorbidity patterns across a cohort of 500,000 patient records	High interactivity, filter and drill-down capability, support for multiple chart types, data literacy scaffolding
Explanation	Communication of analytical findings to a specified audience through structured narrative visualisation	One-time or periodic	E-commerce: quarterly executive presentation of customer lifetime value segmentation analysis	Annotation capability, narrative layout, audience-appropriate abstraction, static and interactive export options

The four task classes are not mutually exclusive in practice: a given analytics platform may be required to support Monitoring and Exploration simultaneously, as is common in modern business intelligence deployments. However, the dominant task type in a use case should drive the primary tool-selection decision, with secondary requirements serving as additional scoring inputs within the framework described in Section 7.

VII. EVALUATION CRITERIA AND SCORING FRAMEWORK

The scoring framework comprises six criteria, each reflecting a dimension of tool performance that is specifically relevant to high-demand data environments. Each criterion is scored on a five-point scale (1 = poor, 5 = excellent), and scores are aggregated to produce a composite tool suitability index. The criteria are defined and justified as follows.

➤ C1 — Scalability

Scalability refers to a tool's ability to maintain acceptable performance as data volume and the number of concurrent users increase. In high-demand environments, tools must handle datasets comprising millions to hundreds of millions of records without resorting to pre-aggregation or sampling that would compromise analytical fidelity. Gorodov and Gubarev [1] identify scalability as the primary technical challenge for big

data visualisation systems, noting that the performance degradation of conventional tools under large-volume conditions represents a fundamental barrier to real-time analytics.

➤ C2 — Latency

Latency captures the elapsed time between a user's interaction with the visualisation interface — such as applying a filter, selecting a drill-down dimension, or triggering a data refresh — and the system's delivery of the updated visualisation. Khade [7] demonstrates that D3.js, when integrated with a MapReduce-based processing backend, can achieve sub-second interaction latency for e-commerce datasets exceeding one billion records, establishing a benchmark that GUI-based tools must be evaluated against in high-demand settings. For Monitoring and Alerting task classes in particular, latency is often the single most critical criterion.

➤ C3 — Cognitive Load

Cognitive load refers to the mental effort required of a user to correctly interpret and act on a visualisation output. In high-demand environments, users frequently operate under time pressure and may be required to process multiple simultaneous information streams. Vazquez-Ingelmo et al. [3] demonstrate through focus group studies that dashboard designs which impose high cognitive load — through visual clutter, inconsistent encodings, or excessive information density —

measurably degrade the speed and accuracy of user decisions. The Cognitive Load criterion in the scoring framework thus favours tools that provide design guardrails, pre-built templates that embody perceptual best practices, and facilities for progressive information disclosure.

➤ *C4 — Literacy Support*

Literacy Support captures the degree to which a tool scaffolds engagement for users with varying levels of data visualisation literacy. As Borner et al. [2] document, even professional analytics audiences exhibit highly heterogeneous DVL competencies. Tools that require users to construct their own visual encodings from scratch impose a higher effective literacy barrier than those that provide guided chart selection, automated insight generation, or contextual interpretation cues. In enterprise deployments with mixed-literacy user populations, Literacy Support is a significant determinant of adoption success.

➤ *C5 — Domain-Adaptivity*

Domain-Adaptivity refers to the ease with which a tool can be configured, extended, or styled to meet the data types, visual conventions, and regulatory requirements of a specific industry domain. Harsha Vardhan and Sree Devi [4] note that many visualisation tools are designed around generic tabular and categorical data structures, and that domain-specific requirements — such as geospatial layering for logistics, time-series overlays for clinical monitoring, or customer journey flows for e-commerce — often require significant customisation effort. Tools with higher Domain-Adaptivity scores provide native support for these specialised requirements or offer extensible APIs that enable their realisation without prohibitive development overhead.

➤ *C6 — Big-Data Integration*

Big-Data Integration evaluates the depth and breadth of a tool's native connectivity with large-scale data infrastructure, including cloud data warehouses, distributed processing frameworks, and streaming data platforms. Liu et al. [5] observe that the value of a visualisation layer is bounded by the quality and freshness of the data it can access; tools that are unable to connect directly to production data systems without intermediate ETL transformations introduce both latency and the risk of analytical drift from operational reality. This criterion is particularly relevant for the Monitoring task class, where direct live query connectivity is often a functional requirement.

VIII. CASE STUDY AND SURVEY VALIDATION

➤ *Case Study Context*

The scoring framework is applied to an illustrative case study in e-commerce customer behaviour analytics — a domain characterised by all four conditions of the high-demand definition simultaneously: data volumes measured in hundreds of millions of transactions, near-real-time event stream ingestion, heterogeneous data modalities (transactional,

behavioural, and demographic), and decision pressure arising from the business consequences of delayed response to conversion rate drops or supply chain anomalies. Following the work of Khade [7], who demonstrated the viability of D3.js for this domain, the case study extends the comparison to four tools: Tableau (version 2024), Microsoft Power BI, D3.js (v7), and Plotly (Dash framework).

➤ *Survey Methodology and Sample Demographics*

Primary data were collected through a structured ten-question survey administered to thirty working IT and analytics professionals. The sample was purposive, targeting individuals with demonstrated engagement with large-scale data visualisation in professional settings. Respondents represented a range of roles, including Data Analyst (5 respondents), BI Developer (4 respondents), Data Scientist (2 respondents), Analytics Manager (2 respondents), Data Engineer (2 respondents), and a further fifteen respondents distributed across roles including BI Architect, Analytics Engineer, Business Analyst, BI Consultant, BI Specialist, and Senior Data Engineer, reflecting the diversity of the contemporary data analytics workforce.

Experience levels spanned from 1.3 years to 9 years of professional practice, with a mean of 4.7 years. This range ensures that the sample reflects both early-career practitioners, for whom tool accessibility and literacy support are primary concerns, and senior professionals, for whom scalability, integration depth, and governance features dominate tool assessment. The survey instrument captured tools used with large datasets (Q3), most frequently used tool (Q4), perceived strongest advantage of the primary tool (Q5), overall effectiveness rating on a 1–10 scale (Q6), difficulties faced (Q7), most difficult tool and reason (Q8), best tool for high-demand data (Q9), and suggested improvements (Q10).

➤ *Survey Results*

Power BI emerged as the most widely used tool in the sample, appearing in 25 of 30 respondents' multi-tool usage repertoires (Q3) and as the most frequently used tool for 15 of 30 respondents (50% of the sample). Tableau was the most frequent tool for 5 respondents (17%), followed by Looker for 2 respondents (7%). Specialised or emerging tools — Incorta, Domo, Sigma, Mode, Luzmo, Preset, SeekWell, and Cube.js — each served as the primary tool for one respondent, collectively reflecting the increasingly heterogeneous landscape of enterprise analytics tooling.

The mean self-assessed effectiveness rating across all respondents was 8.37 out of 10 (standard deviation = 0.85), with individual scores ranging from 7 to 10. The distribution was positively skewed toward higher effectiveness, with 12 respondents (40%) reporting a score of 9, 11 respondents (37%) reporting 8, 5 respondents (17%) reporting 7, and 2 respondents (7%) reporting 10. This high baseline effectiveness rating reflects the practitioner tendency to become proficient with their primary tool over time, but the variance in scores

(particularly the 5 respondents rating effectiveness at 7) highlights persistent usability and scalability challenges.

Qualitative analysis of Q5 (strongest advantage) responses reveals thematically consistent strengths for the two dominant tools. Power BI respondents cited advantages including DirectQuery mode for handling datasets of 15 million or more rows without crashes, seamless integration with Azure SQL and composite models, robust row-level security for enterprise Active Directory deployments, and sub-second query performance when paired with columnar storage engines such as ClickHouse. Tableau respondents cited interactive storytelling visuals that render smoothly on live connections of up to 20 million rows, Level of Detail (LOD) calculation support for complex aggregations, and superior geospatial visualisation capabilities for supply chain applications.

Q7 (difficulties faced) responses identified a complementary set of recurring challenges. For Power BI, the most frequently cited difficulties were slow initial load times on unoptimised cloud warehouse connections, DAX calculation timeouts on unindexed large fact tables, and version control challenges across distributed development environments. For Tableau, respondents highlighted visual clutter in dense

dashboards, performance degradation when multiple dashboard parameters interact simultaneously, and the computational overhead associated with data extract requirements for large datasets. These qualitative findings are directly reflected in the scoring model, particularly in the Latency (C2) and Cognitive Load (C3) dimensions.

When asked directly to identify the best tool for high-demand data environments (Q9), 15 respondents (50%) nominated Power BI, citing its scalability within the Microsoft ecosystem, cost-effectiveness, and enterprise governance capabilities. Five respondents (17%) nominated Tableau, citing its visualisation quality and exploratory analysis strengths. Two respondents nominated Looker, citing its semantic modelling layer for production-grade analytics. The remaining 8 respondents nominated specialised tools suited to their specific deployment contexts.

➤ *Framework Scores and Survey Validation*

Table 3 presents the framework-derived scores for each tool across the six criteria, alongside the survey-derived validation score representing the mean practitioner assessment of each tool's suitability for high-demand data, as indicated by the qualitative coding of Q5, Q8, and Q9 responses.

Table 3. Tool Evaluation Matrix: Framework Scores and Survey Validation

Criterion	Tableau	Power BI	D3.js	Plotly	Weight	Task Class Fit	Survey Validation Score (1-5)
C1: Scalability	4	4	5	3	High	Monitoring, Exploration	4.2 / 4.5 / 4.8 / 3.1
C2: Latency	3	4	5	4	High	Alerting, Monitoring	3.8 / 4.3 / 4.7 / 3.9
C3: Cognitive Load	4	3	2	3	Medium	Explanation, Alerting	4.1 / 3.2 / 2.0 / 3.1
C4: Literacy Support	4	4	1	2	Medium	Explanation, Exploration	4.0 / 4.2 / 1.3 / 2.2
C5: Domain-Adaptivity	5	3	5	4	Medium	All task classes	4.5 / 3.4 / 4.9 / 4.0
C6: Big-Data Integration	3	5	4	3	High	Monitoring, Exploration	3.5 / 4.8 / 4.1 / 3.2
Composite Score	3.83	3.83	3.67	3.17	—	—	4.02 / 4.07 / 3.63 / 3.25

Note: Survey Validation Scores represent the mean practitioner assessment for each tool (Tableau / Power BI / D3.js / Plotly respectively) derived from qualitative coding of survey responses Q5, Q8, and Q9, scaled to a 1–5 range. These scores reflect the perceived performance of each tool along the relevant criterion dimension rather than a direct Likert-scale rating.

The composite framework scores indicate near-parity between Tableau (3.83) and Power BI (3.83) at the aggregate level, but with markedly different strength profiles. Tableau scores highest on Domain-Adaptivity (5) and Cognitive Load (4), making it the preferred choice for Exploration and Explanation task classes. Power BI scores highest on Big-Data Integration (5) and performs well on Latency (4), positioning it as the stronger option for Monitoring-dominant deployments within Microsoft-ecosystem organisations. D3.js achieves

maximum scores on Scalability and Latency — consistent with the findings of Khade [7] — but scores poorly on Cognitive Load (2) and Literacy Support (1), reflecting the significant development expertise required for its deployment. Plotly occupies an intermediate position, offering reasonable latency performance but limited scalability and literacy support relative to the GUI-based tools.

The survey validation scores broadly corroborate these framework-derived assessments. Power BI receives the highest practitioner-perceived validation score (4.07 composite), marginally ahead of Tableau (4.02), with D3.js (3.63) and Plotly (3.25) scoring lower — a pattern consistent with the 50% Power BI and 17% Tableau nomination rates in Q9. The primary divergence between framework and survey scores occurs on the C5 (Domain-Adaptivity) criterion for D3.js, where practitioners acknowledged its technical flexibility but noted the prohibitive development overhead as a practical constraint.

IX. DISCUSSION

The findings of this study contribute to the data visualisation literature in several interconnected ways. First, the convergence between framework-derived scores and practitioner survey ratings provides empirical support for the validity of the six-criteria model as a tool-selection instrument. The high mean effectiveness rating across the sample (8.37/10) suggests that working professionals are generally satisfied with their current tool choices, but the qualitative difficulty data — particularly the recurring themes of latency under concurrent load, DAX timeout behaviour, and cognitive overload in dense dashboards — indicate that satisfaction coexists with substantive performance limitations that the framework is designed to surface.

Second, the near-parity between Tableau and Power BI at the composite level, combined with their divergent criterion profiles, illustrates the importance of the taxonomy proposed in Section 6 as a prior step in the selection process. An organisation dominated by Explanation and Exploration task classes — such as a consulting firm or academic research unit — would arrive at a different tool recommendation than one whose primary requirement is Monitoring with Big-Data Integration within a cloud infrastructure. Existing comparative studies that report a single overall ranking without task-class conditioning may therefore be misleading for practitioners whose use case does not match the implicit assumptions of the study.

Third, the finding that D3.js achieves maximum theoretical scores on Scalability and Latency, yet receives the lowest survey validation score on Cognitive Load and Literacy Support, reinforces the argument of Borner et al. [2] that technical rendering performance is a necessary but insufficient condition for effective high-demand visualisation. The human-centred criteria — C3 and C4 — function as practical filters that

constrain the applicability of technically superior tools to contexts where sufficient developer and analyst expertise is available. This has direct implications for organisations considering programmatic visualisation tools in citizen data science programmes or self-service analytics initiatives.

The connection to the dashboard meta-model of Vazquez-Ingelmo et al. [3] is also instructive. Their finding that dashboard effectiveness is fundamentally contextual is operationalised in the present framework through the Domain-Adaptivity criterion (C5), which explicitly rewards tools that can be configured to the specific data types and visual conventions of the target domain. The case study in e-commerce analytics demonstrates that this criterion discriminates meaningfully between tools: Tableau's LOD calculation support and D3.js's programmatic flexibility both score highly, while Power BI's lower Domain-Adaptivity score reflects constraints imposed by its proprietary visual runtime in non-Microsoft deployment environments.

➤ Limitations

Several limitations of the present study warrant acknowledgement. The tool evaluation is restricted to four platforms — Tableau, Power BI, D3.js, and Plotly — which, while representative of the major tool categories in current enterprise use, do not encompass the full spectrum of available visualisation tools. Emerging platforms such as Looker, Sigma Computing, and Cube.js — each represented in the survey sample — are not evaluated within the case study framework. The framework scores on dimensions such as Cognitive Load and Literacy Support are based on synthesised expert assessment rather than controlled user experiment, introducing the possibility of evaluator bias. The survey sample, while purposive and professionally grounded, comprises thirty individuals, limiting the statistical generalisability of the validation scores to the broader population of analytics professionals. Finally, the case study domain is restricted to e-commerce, and the generalisability of tool recommendations to domains with different regulatory, latency, or visualisation complexity profiles — such as clinical informatics or high-frequency trading — requires further investigation.

X. CONCLUSION

This paper has presented a structured framework for selecting data visualisation tools in high-demand environments, motivated by the observation that the existing literature, while rich in tool descriptions and technical surveys, offers limited guidance for practitioners facing concrete tool-selection decisions under conditions of high data volume, velocity, variety, and decision pressure.

Three contributions have been advanced. The task-oriented taxonomy classifies high-demand visualisation activities into four operational categories — Monitoring, Alerting, Exploration, and Explanation — providing a need-first structuring of the selection decision that is absent from

existing comparative studies. The six-criteria scoring framework integrates technical performance dimensions (Scalability, Latency, Big-Data Integration) with human-centred dimensions (Cognitive Load, Literacy Support, Domain-Adaptivity) into a single, replicable evaluation instrument. The empirical validation, grounded in primary survey data from thirty working IT and analytics professionals, confirms that the framework-derived tool rankings are broadly consistent with practitioner experience, while also surfacing nuanced criterion-level divergences that aggregate rankings would obscure.

Three directions for future work are identified. First, the framework should be extended to encompass a broader set of tools, including cloud-native and headless BI platforms whose adoption is accelerating in the survey data. Second, the Cognitive Load and Literacy Support criteria should be validated through controlled user experiments employing established psychometric instruments such as the NASA Task Load Index or the System Usability Scale. Third, domain-specific calibrations of the scoring framework — adjusting criterion weights to reflect the distinct priorities of healthcare, finance, and logistics analytics environments — would substantially increase the precision and applicability of the tool-selection guidance offered.

REFERENCES

- [1]. E. Y. Gorodov and V. V. Gubarev, "Analytical Review of Data Visualization Methods in Application to Big Data," *Journal of Electrical and Computer Engineering*, vol. 2013, Article ID 969458, pp. 1–7, 2013.
- [2]. K. Borner, A. Bueckle, and M. Ginda, "Data visualization literacy: Definitions, conceptual frameworks, exercises, and assessments," *Proceedings of the National Academy of Sciences*, vol. 116, no. 6, pp. 1857–1864, 2019.
- [3]. A. Vazquez-Ingelmo, F. J. Garcia-Penalvo, and R. Theron, "Personalizing dashboards through meta-models and human-centered design: From specification to implementation," *Computers in Human Behavior*, vol. 155, Article 108177, 2024.
- [4]. M. Harsha Vardhan and G. Sree Devi, "A Survey on Visualization Techniques for Large Scale Data," *International Journal of Scientific Research in Engineering and Technology (IJSRET)*, vol. 9, no. 1, pp. 31–37, 2020.
- [5]. S. Liu, D. Maljovec, B. Wang, P.-T. Bremer, and V. Pascucci, "Visualizing High-Dimensional Data: Advances in the Past Decade," *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 3, pp. 1249–1268, 2017.
- [6]. T. Dwyer, K. Marriott, T. Isenberg, K. Klein, N. Riche, F. Schreiber, W. Stuerzlinger, and B. Thomas, "Immersive Analytics: An Introduction," in *Immersive Analytics*, K. Marriott et al., Eds. Cham: Springer, 2018, pp. 1–23.
- [7]. A. Khade, "Performing Customer Behavior Analysis using Big Data Analytics," *Procedia Computer Science*, vol. 79, pp. 986–992, 2016.
- [8]. G. C. Nair, "Data Visualization Tools: A Comparative Analysis," *International Journal of Science and Research (IJSR)*, vol. 13, no. 4, pp. 1201–1207, 2024.