




International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: <https://www.svedbergopen.com/>



Research Paper

Open Access

Unified Semantic Layers: Eliminating Metric Inconsistency Across Distributed Enterprise Analytics

Ashwin Krishnappa Kumar

Senior Manager of Data Engineering and Analytics, Tesla, USA. ORCID: 0009-0007-7020-4227

Abstract

Distributed enterprise analytics has become the default operating condition for large organizations rather than an exception, as firms decentralize data ownership across business units, regional operations, and product lines. A recurring consequence of this decentralization is that the same underlying business concept, such as revenue, active users, or churn, is routinely defined, calculated, and reported differently by different teams querying different systems, even when the row-level data feeding those calculations is fully accurate. This paper argues that this inconsistency is best understood as a governance and semantics defect rather than a traditional data quality defect and that a unified semantic layer, a governed, queryable abstraction that encodes metric definitions, business logic, and access rules in one place rather than in each downstream tool separately, is the systems-engineering mechanism that resolves it. The paper synthesizes foundational data quality and query-view theory, the empirical data mesh and data governance literature, and recent practitioner-facing research on semantic layers and business intelligence adoption into a single framework of four design principles that a unified semantic layer must satisfy: architecturally enforced single definitional authority, domain-federated authorship with centrally governed publication, view-based logical abstraction over physically distributed sources, and consumer-agnostic accountability extending to both human and AI-mediated consumers. The synthesis further shows that the technical mechanism alone is insufficient without organizational conditions long identified in the literature on business intelligence adoption, including perceived trustworthiness of system output and baseline governance maturity. This paper is presented explicitly as an integrative literature synthesis grounded in verified published findings rather than an original instrumented deployment study. The analysis concludes that metric inconsistency in decentralized enterprise architectures is fundamentally a governance problem implemented through a technical mechanism and that the semantic layer's distinct contribution is converting metric governance from a voluntary organizational practice into an architectural constraint.

Keywords: Semantic Layer, Data Mesh, Metric Consistency, Data Governance, Distributed Enterprise Analytics, Data Virtualization

1. Introduction

Distributed enterprise analytics has become the default operating condition for large organizations rather than an exception. As firms decentralize data ownership across business units, regional operations, and product lines, the same underlying business concept, revenue, active user, order, and churn are routinely defined, calculated, and reported differently by different teams querying different systems. A finance team's definition of monthly active revenue may exclude refunds still in adjudication while a product analytics team's definition of the same metric includes them, producing two authoritative-looking numbers for what should be one fact. This is not a data quality defect in the traditional sense of missing or malformed values. It is a governance and semantics defect: the data itself may be perfectly accurate at the row level while the aggregate metric built from it diverges silently across consumers.

This inconsistency carries a measurable organizational cost. When two dashboards report different figures for what business stakeholders believe is the same metric, the practical effect is not merely an inconvenience but an erosion of trust in analytics as a decision-making input, prompting manual reconciliation work that itself introduces further definitional drift [1]. The problem has intensified as centralized data warehouse architectures have given way to more decentralized patterns, including domain-oriented data mesh architectures that deliberately distribute data ownership across autonomous teams as a scaling strategy [2]. Decentralization solves a genuine bottleneck problem: a single centralized data team cannot indefinitely absorb the modeling and pipeline demands of an entire enterprise, but it reintroduces, at organizational scale, exactly the metric consistency problem that a single centralized warehouse implicitly avoided by construction.

This paper argues that a unified semantic layer, a governed, queryable abstraction that sits between raw distributed data sources and the tools that consume them and that encodes metric definitions, business logic, and access rules in one place rather than in each downstream tool separately, is the systems-engineering mechanism that resolves this tension. The semantic layer allows an organization to keep the operational

benefits of distributed data ownership, and domain teams retain control of their data products while restoring the single-definition guarantee that centralized warehouses provided implicitly. This is not a new observation in isolation; database theory has studied the general problem of answering queries against a shared logical view for decades [3], and the modern semantic layer is, in an important sense, a direct architectural descendant of that query-rewriting and view-based literature applied to a modern, enterprise-scale problem.

The contribution of this paper is a synthesis, drawn from foundational data quality and query-view theory, from the empirical data mesh and data governance literature, and from recent practitioner-facing research on semantic layers, into a single framework that describes the design principles a unified semantic layer must satisfy to actually eliminate metric inconsistency rather than merely relocate it. The paper is explicitly framed as an integrative literature synthesis grounded in verified published findings rather than an original instrumented deployment study, consistent with the honest scope of the analysis that follows.

2. The Metric Inconsistency Problem in Distributed Enterprise Analytics

The foundational literature on data quality distinguishes accuracy, whether a value correctly represents the real-world fact it claims to represent, from a broader set of quality dimensions that are consumer-relative rather than absolute, including consistency, interpretability, and believability [1]. This distinction matters directly for the metric inconsistency problem. A revenue figure can be perfectly accurate at the transactional row level, every individual sale correctly recorded, and still produce inconsistent aggregate metrics if two consuming teams apply different filtering, windowing, or currency-conversion logic when rolling those rows up into a single number. Wang and Strong's framework establishes that data quality is not a property of the data alone but of the data as interpreted and consumed, which is precisely the property that breaks down when the same underlying data is aggregated through multiple, independently maintained business logic layers [1]. The scale of this problem has grown in step with the decentralization of enterprise data architecture itself. Recent empirical research on data-driven decision-making finds that analytical culture and the degree of centralization in an organization's data function are themselves significant predictors of whether data-driven decisions are actually adopted in practice, independent of the sophistication of the analytics tooling deployed [4]. This finding is consistent with the practical experience that inconsistent metrics are frequently not a tooling failure but an organizational and governance failure: multiple teams independently encode the same business logic with small, undocumented variations, and no single system of record adjudicates between them.

Case study research on data governance implementation reinforces this point directly. In a study spanning twenty organizational case studies, Walsh, McAvoy, and Sammon find that formal data governance theory frequently does not resonate with practitioners in the way academic frameworks assume, and that the human and behavioral dimensions of managing how people actually define, request, and modify data logic have a larger effect on program success than the specific governance framework adopted [5]. This evidence is a significant finding for the semantic layer argument developed in this paper: a semantic layer is not merely a technical artifact but a governance artifact, and its success depends on whether the organization treats the metric definitions it encodes as the single, socially agreed source of truth rather than as one more optional reference that teams may still bypass.

A parallel empirical study of data governance in digital ecosystems, examining multiple holistic case studies of organizations operating across ecosystem boundaries rather than within a single firm, finds that governance challenges compound further when data crosses organizational boundaries, since ecosystem partners frequently lack the direct authority relationships that make governance enforceable within a single firm [6]. While this paper's scope is the distributed enterprise rather than the multi-firm ecosystem, the underlying mechanism, that governance authority and data ownership are frequently misaligned once data is decentralized, is directly analogous and strengthens the case that a technical mechanism alone, without a clear governance and ownership model, cannot resolve metric inconsistency.

3. From Data Warehouses to Data Mesh: The Evolution of Distributed Analytics Architecture

Centralized data warehousing achieved metric consistency largely by construction rather than by explicit governance design. When a single centralized team owns the extraction, transformation, and loading pipeline for the entire enterprise, and every reporting tool queries the same warehouse schema, there is structurally only one place where a metric definition can be encoded, so inconsistency across consumers is architecturally difficult to produce even without an explicit consistency mechanism. The tradeoff of this architecture is well documented in more recent literature examining why organizations have moved away from it: a single centralized data team becomes a bottleneck as the number of data-producing domains and data-consuming use cases grows, since every new source system integration and every new metric request routes through the same limited team [2].

Data mesh architecture, as formalized by Machado, Costa, and Santos, responds to this bottleneck by decentralizing data ownership to domain teams, treating data as a product with a defined owner, and by introducing a self-serve platform for data infrastructure that allows domain teams to publish and consume data products without depending on a central pipeline team for every change [2]. A systematic gray literature review of data mesh implementations across industry sources by Goedegebuure et al. confirms that this decentralization genuinely resolves the organizational scaling bottleneck that motivated it, while simultaneously surfacing a well-documented and recurring secondary problem: without a strong federated governance layer, decentralized domain teams tend to produce data products with inconsistent modeling conventions, inconsistent naming, and, most relevant to this paper, inconsistent metric definitions across domains [7].

This issue is not an incidental implementation detail but a structural consequence of the architecture. Data mesh explicitly trades the single central chokepoint that made warehouse-era consistency easy for the operational flexibility of domain autonomy, and the literature is consistent in identifying federated computational governance, a governance layer that domain teams comply with programmatically rather than through voluntary policy adherence, as the necessary counterbalancing mechanism [2], [7]. Without it, the same metric inconsistency problem that motivated centralized warehousing in the first place, reappears, now distributed across as many independently evolving definitions as there are domain teams.

Gieß and Hutterer provide a broader academic delimitation of the current distributed data architecture landscape, situating data mesh alongside related but distinct patterns, data platforms, data spaces, and data fabrics, and clarifying that these architectures are frequently conflated in practitioner discourse despite addressing different combinations of the ownership-decentralization and technical-integration problems [8]. Their analysis is useful for this paper because it makes explicit that decentralized ownership and technical integration are separable concerns: an organization can decentralize data ownership without also decentralizing the semantic definitions that make the resulting data products comparable and combinable, which is precisely the design space a unified semantic layer occupies.

4. The Semantic Layer as a Governance and Consistency Mechanism

The technical foundation for a semantic layer predates the current data mesh discourse by decades. The general problem of answering a query against a logical, potentially virtual view rather than directly against physical base tables, allowing the underlying physical storage to change or be distributed while the logical query surface remains stable, was formalized in the database theory literature by Halevy, whose survey of query-answering-using-views techniques established the theoretical machinery, query containment, view-based rewriting, and completeness conditions under which a view-based answer is guaranteed equivalent to a direct query that a modern semantic layer relies on when it presents a single governed metric definition over data that may physically reside across multiple distributed domain-owned sources [3]. A semantic layer is, in this precise technical sense, an applied instance of the view-based query answering problem, specialized to business metrics rather than arbitrary relational queries.

Recent practitioner-oriented research from MIT's Center for Information Systems Research makes the organizational case for this technical mechanism directly. Lefebvre, Wixom, Legner, van der Meulen, and Beath argue that the semantic layer's central value is not primarily performance or convenience but trust: by encoding metric logic once, in a governed layer that every downstream tool queries rather than reimplements, organizations create a single point of accountability for what a metric means, which in turn allows business stakeholders to trust a number without needing to independently audit the query logic behind every dashboard that produces it [9]. This framing directly supports the argument advanced in this paper, that the semantic layer's contribution is best understood as a governance mechanism implemented in technical form, rather than as a performance optimization that happens to also improve consistency.

This governance framing is reinforced by systematic research on data governance frameworks more broadly. Acev et al., in a systematic analysis of data governance frameworks and their applicability to environments involving shared or federated data assets, find that the frameworks best suited to federated and multi-party data environments consistently share a common structural feature: a clearly designated, technically enforced point of definitional authority, rather than a purely policy-based or voluntary compliance model [10]. A semantic layer operationalizes exactly this structural feature for metric definitions specifically: rather than asking every consuming team to voluntarily comply with a shared metric glossary maintained in a separate document, the semantic layer makes the governed definition the only technically available path to querying that metric, converting governance from a voluntary practice into an architectural constraint.

The same governance logic extends into adjacent domains. Desai and Desai, examining governance frameworks for large language model outputs used in enterprise analytics contexts, identify an analogous pattern: as

organizations increasingly use generative AI systems to query and summarize analytics data in natural language, the risk of the AI system inventing or inconsistently interpreting metric definitions grows unless there is a single governed semantic source the AI system is constrained to query against rather than a set of independently maintained dashboards it might paraphrase inconsistently [11]. This finding extends the urgency of the unified semantic layer argument beyond human dashboard consumers to an emerging class of AI-mediated analytics consumers, for whom an ungoverned, inconsistent metric landscape poses an even sharper risk, since an AI system cannot exercise the same contextual judgment a human analyst might apply when reconciling two conflicting numbers.

5. Synthesized Framework: Design Principles for a Unified Semantic Layer

Drawing together the foundational query-view theory, the literature on data mesh governance, and the recent semantic layer research reviewed above, this synthesis proposes four design principles that a unified semantic layer must satisfy to eliminate metric inconsistency across a distributed enterprise, rather than merely adding one more optional reporting layer alongside the existing inconsistent ones.

The first principle is single definitional authority, enforced architecturally rather than procedurally. Consistent with Acev et al.'s finding that effective federated governance frameworks rely on technically enforced rather than voluntary compliance points [10], a unified semantic layer must be the only sanctioned path by which any downstream tool retrieves a governed metric. If dashboards and reporting tools retain the technical ability to bypass the semantic layer and query underlying tables directly, the governance guarantee collapses regardless of how well the semantic layer itself is designed.

The second principle is domain-federated authorship with centrally governed publication. This principle directly reconciles the decentralization benefits of data mesh architecture [2] with the consistency benefits of centralized warehousing that motivated this paper's argument. Domain teams should retain authorship of the metric logic relevant to their own domain, since they hold the necessary business context, but that logic should pass through a central, lightweight governance and publication step before becoming queryable enterprise-wide, mirroring the federated computational governance pattern identified as necessary in the data mesh gray literature review [7].

The third principle is view-based logical abstraction over physically distributed sources, grounded directly in Halevy's query-answering-using-views formalization [3]. The semantic layer's metric definitions should be expressed as logical views with well-defined containment and equivalence properties relative to the underlying distributed domain data products, allowing the physical location and ownership of the underlying data to change, consistent with ongoing data mesh domain evolution, without invalidating the metric definition or requiring every downstream consumer to be updated.

The fourth principle is consumer-agnostic accountability, extending the trust argument advanced by Lefebvre et al. [9] to explicitly include AI-mediated consumers alongside human dashboard users, consistent with the governance risk identified by Desai and Desai for generative AI systems operating over enterprise analytics data [11]. The semantic layer's governed definitions should be the mandatory query surface for any consumer of enterprise metrics, human or automated, rather than a resource that only traditional business intelligence tools are expected to respect. Table 1 summarizes these four principles alongside the literature that grounds each one.

Design Principle	Description	Grounding Literature
1. Single Definitional Authority	Enforced architecturally, not procedurally, the semantic layer is the only sanctioned path to a governed metric	[10]
2. Domain-Federated Authorship	Domain teams author metric logic for their own domain; a central, lightweight step governs publication enterprise-wide	[2],[7]
3. View-Based Logical Abstraction	Metric definitions expressed as logical views over physically distributed domain sources, with containment/equivalence guarantees	[3],[12]
4. Consumer-Agnostic Accountability	Governed definitions are the mandatory query surface for both human and AI-mediated consumers	[9],[11]

Table 1. The four design principles and their grounding literature.

5.1 Technical Implementation: Data Virtualization as the Substrate for View-Based Abstraction

The third design principle, view-based logical abstraction over physically distributed sources, requires a concrete technical substrate capable of executing federated queries against domain-owned sources without

first physically consolidating that data into a single warehouse. Data virtualization architecture provides this substrate. Balakrishnan's framework for multi-engine data virtualization describes an architecture in which a virtualization layer maintains logical schema mappings to multiple underlying physical engines, relational databases, object stores, and streaming systems among them, and rewrites incoming queries against the logical schema into engine-specific execution plans at query time [12]. This approach is a direct practical realization of the view-based query answering machinery formalized decades earlier [3], applied specifically to the multi-engine, multi-domain reality of a modern distributed enterprise rather than to a single relational database. The practical significance for the unified semantic layer argument is that the semantic layer's governed metric definitions need not require domain teams to migrate their underlying data into a shared physical store, which would reintroduce exactly the central-bottleneck problem that motivated the move to data mesh architecture in the first place [2]. A data-virtualization-backed semantic layer instead leaves domain-owned data physically where it already resides and enforces the single-definition guarantee purely at the logical query layer, satisfying both the domain autonomy goal of data mesh and the definitional consistency goal that this paper has argued is otherwise put at risk by that same decentralization.




Figure 1. Unified semantic layer architecture: domain sources, data virtualization/view layer, governed semantic layer, and human and AI-mediated consumers.

6. Discussion: Organizational and Technical Implications

Table 2 compares the centralized warehouse, data mesh, and unified semantic layer approaches across three dimensions central to this paper's argument: data ownership, the mechanism relied on for metric consistency, and the governance model in force.

Dimension	Centralized Warehouse	Data Mesh	Unified Semantic Layer
Data ownership	Single central data team	Decentralized, domain-owned data products	Domain-owned data, centrally governed definitions
Consistency mechanism	Implicit, by construction (one pipeline, one schema)	Not guaranteed by architecture alone	Explicit, architecturally enforced single definitional authority

Governance model	Centralized, procedural	Federated, often voluntary/policy-based [7]	Federated computational governance, technically enforced [10]
------------------	-------------------------	---	---

Table 2. Comparative summary: centralized warehouse, data mesh, and unified semantic layer.

Two implications follow directly from this synthesis. First, the metric inconsistency problem in distributed enterprises is not primarily solvable through better data quality tooling in the traditional sense of validation rules and anomaly detection, since, as Wang and Strong's framework establishes, the underlying data can be fully accurate while still producing inconsistent aggregate metrics [1]. Organizations that invest exclusively in data quality monitoring without addressing the semantic layer gap should expect metric inconsistency to persist, because the defect these tools are designed to catch is not the defect that is actually producing the inconsistency.

Second, the human and organizational dimension identified in case study research on governance program success [5] means that a technically well-designed semantic layer can still fail to eliminate metric inconsistency if the organization does not also address the behavioral incentive for teams to bypass it, for example when a domain team's urgent reporting deadline makes querying the underlying table directly faster than requesting a new metric definition through the governed publication process. This paper's synthesis therefore frames the semantic layer as a necessary but not independently sufficient condition: it must be paired with an organizational commitment, reflected in the finding that analytical culture and centralization decisions independently predict data-driven decision adoption [4], to treat the governed definition as genuinely authoritative rather than as an optional convenience.

A further implication concerns architectural sequencing for organizations currently mid-migration from centralized warehousing toward data mesh patterns. Because the metric inconsistency risk identified in this paper arises specifically from the decentralization step [2], [7], the synthesis suggests that the semantic layer's federated governance mechanism should be established before or concurrently with domain decentralization, rather than retrofitted afterward once inconsistent domain-specific metric definitions have already proliferated and become embedded in downstream reporting habits.

A unified semantic layer, however well designed technically, is ultimately one more enterprise system that must be adopted by the domain teams and business consumers it is meant to serve, and the literature on business intelligence adoption offers directly relevant findings on what determines whether such adoption succeeds. Ain, Vaia, DeLone, and Waheed's systematic review of two decades of research on business intelligence systems identifies organizational readiness, senior management sponsorship, and, notably, the perceived trustworthiness of the system's output as recurring, empirically supported predictors of sustained BI system adoption and utilization, findings drawn from a substantial body of prior empirical studies rather than a single case [13]. Trustworthiness of output maps directly onto this paper's central argument: a semantic layer that consumers do not yet trust as authoritative will not displace the ad hoc, inconsistent metric definitions it is meant to replace, regardless of its technical soundness, since users will continue to fall back on familiar, if inconsistent, calculation habits until the governed layer earns that trust.

A more recent systematic literature review by Horani, Khatibi, AL-Soud, Tham, and Al-Adwan, focused specifically on organizational-level business analytics adoption factors, extends this finding by identifying data governance maturity itself as a significant antecedent of successful analytics adoption, rather than a parallel, independent initiative [14]. This is a meaningful reinforcement of the sequencing implication raised earlier in this paper's discussion: organizations attempting to roll out a unified semantic layer without first establishing baseline data governance maturity, clear data ownership, agreed escalation paths for definitional disputes, are working against a documented adoption headwind rather than a neutral starting condition.

Joshi's applied research on integrating data governance with advanced analytics to improve enterprise decision-making provides a closing empirical link between the governance mechanism this paper has centered and the decision-making outcomes that ultimately justify investing in it, finding that organizations with integrated governance-analytics programs report measurably faster and more confident decision cycles than organizations that treat governance and analytics capability as separate initiatives [15]. This finding supports the paper's framing of the unified semantic layer not as a defensive, compliance-oriented control but as a direct enabler of the faster, more confident data-driven decision-making that centralization-and-culture research has already shown depends on organizational, not merely technical, conditions [4].

7. Conclusion

This paper has synthesized foundational data quality and query-view theory, the empirical data mesh and data governance literature, and recent semantic layer research into a single framework of four design principles,

architecturally enforced single definitional authority, domain-federated authorship with centrally governed publication, view-based logical abstraction over physically distributed sources, and consumer-agnostic accountability extending to AI-mediated consumers, that together describe how a unified semantic layer eliminates metric inconsistency across distributed enterprise analytics. The central argument is that metric inconsistency in decentralized architectures is a governance problem implemented through a technical mechanism, not a data quality problem in the traditional sense, and that the semantic layer's value lies specifically in converting metric governance from a voluntary organizational practice into an architectural constraint. This synthesis is presented as an integrative literature review grounded in verified published findings rather than as an original instrumented deployment study, and this framing should be read as an honest limitation rather than a diminishment of its contribution. Future work extending this synthesis with instrumented case study data on semantic layer adoption outcomes, particularly measuring the reduction in metric-reconciliation labor before and after adoption, would substantially strengthen the empirical grounding of the framework proposed here.

Author Contributions

Ashwin Kumar is the sole author of this manuscript and is responsible for conceptualization, literature synthesis, analysis, and writing.

Funding

This research received no external funding.

Conflicts of Interest

The author declares no conflicts of interest.

Data Availability Statement

Not applicable. This article is a synthesis of previously published sources; no new data were generated.

References

- [1] R. Y. Wang and D. M. Strong, "Beyond Accuracy: What Data Quality Means to Data Consumers", *Journal of Management Information Systems*, 1996. Available: <https://www.tandfonline.com/doi/abs/10.1080/07421222.1996.11518099>
- [2] I. A. Machado, C. Costa and M. Y. Santos, "Data Mesh: Concepts and Principles of a Paradigm Shift in Data Architectures", *Procedia Computer Science*, 2022. Available: <https://doi.org/10.1016/j.procs.2021.12.013>
- [3] A. Y. Halevy, "Answering Queries Using Views: A Survey", *The VLDB Journal*, 2001. Available: <https://doi.org/10.1007/s007780100054>
- [4] B. Szukits and G. Móricz, "Towards Data-Driven Decision Making: The Role of Analytical Culture and Centralization Efforts", *Review of Managerial Science*, 2024. Available: <https://doi.org/10.1007/s11846-023-00694-1>
- [5] F. Volz, C. Münch, M. Lohmüller and C. Küffner, "From Data Jungle to Data Governance in Digital Ecosystems: Empirical Evidence From a Multiple Holistic Case Study", *Journal of Business Research*, 2025. Available: <https://doi.org/10.1016/j.jbusres.2025.115747>
- [6] M. J. Walsh, J. McAvoy and D. Sammon, "The Data Governance Journey in Practice: Insights From Case Study Research", *Information Systems Management*, 2025. Available: <https://doi.org/10.1080/10580530.2025.2477459>
- [7] A. Goedegebuure et al., "Data Mesh: A Systematic Gray Literature Review", *ACM Computing Surveys*, 2024. Available: <https://doi.org/10.1145/3687301>
- [8] A. Gieß and A. Hutterer, "The Future of Data Management: A Delimitation of Data Platforms, Data Spaces, Data Meshes, and Data Fabrics", *Information Systems and e-Business Management*, 2025. Available: <https://doi.org/10.1007/s10257-025-00707-4>
- [9] H. Lefebvre, B. H. Wixom, C. Legner, N. van der Meulen and C. M. Beath, "The Case for a Semantic Layer", MIT Center for Information Systems Research Research Briefing, 2026. Available: https://cistr.mit.edu/publication/2026_0501_SemanticLayer_LefebvreWixomLegnerVandermeulenBeath
- [10] D. Acev, S. Biyani, F. Rieder, T. T. Aldenhoff, M. Blazevic, D. M. Riehle and M. A. Wimmer, "Systematic Analysis of Data Governance Frameworks and Their Relevance to Data Trusts", *Management Review Quarterly*, 2025. Available: <https://doi.org/10.1007/s11301-025-00545-1>

- [11] D. Desai and A. Desai, "From Models to Metrics: A Governance Framework for Large Language Models in Enterprise AI and Analytics", *Analytics*, 2026. Available: <https://doi.org/10.3390/analytics5010008>
- [12] A. Balakrishnan, "Data Virtualization Architecture Framework Using Multi-Engine Data Platforms for Big Data Analytics and Machine Learning", *International Journal of Computer Trends and Technology*, 2024. Available: <https://doi.org/10.14445/22312803/IJCTT-V72I2P115>
- [13] N. Ain, G. Vaia, W. H. DeLone and M. Waheed, "Two Decades of Research on Business Intelligence System Adoption, Utilization and Success - A Systematic Literature Review", *Decision Support Systems*, 2019. Available: <https://doi.org/10.1016/j.dss.2019.113113>
- [14] O. M. Horani, A. Khatibi, A. R. AL-Soud, J. Tham and A. S. Al-Adwan, "Determining the Factors Influencing Business Analytics Adoption at Organizational Level: A Systematic Literature Review", *Big Data and Cognitive Computing*, 2023. Available: <https://doi.org/10.3390/bdcc7030125>
- [15] D. Joshi, "Integrating Data Governance and Advanced Analytics to Improve Enterprise Decision-Making", *International Journal of Computational and Experimental Science and Engineering*, 2023. Available: <https://doi.org/10.22399/ijcesen.4512>