ISSN: 2473-4901
v12 n6678

Bridging the Gap: Transitioning Business Intelligence Curricula from On-Premise Data Warehousing to Cloud Analytics

Wenqi Zhou
Zhouw@duq.edu

Sara Khan
khans10@duq.edu

Duquesne University
Pittsburgh, PA, 15237, U.S.

Abstract

Modern Business Intelligence (BI) is undergoing a significant architectural shift from isolated, on-premise data warehousing to agile, cloud-based environments capable of supporting AI-augmented analytics. Despite this industry mandate, undergraduate BI curricula often struggle to bridge the gap between teaching foundational data modeling and modern cloud federation due to high software barriers and a lack of unified textbooks. This pedagogical case study presents a replicable, 15-week curriculum framework that successfully transitions senior-level undergraduate students through the architectural evolution of BI. Utilizing experiential learning resources from the SAP University Alliance, the course begins with strict, on-premise dimensional modeling (SAP BW4/HANA), bridges to cloud-based data federation (SAP Datasphere), and culminates in predictive data mining and AI-augmented dashboarding (SAP Analytics Cloud). By maintaining a consistent focus on structured data throughout the pipeline, students experience a controlled architectural comparison, allowing them to fully grasp the integration values of cloud computing. This paper provides a practical roadmap for Information Systems educators seeking to modernize their BI offerings, demonstrating how to equip students with the adaptable, hybrid-cloud competencies demanded by the modern data-driven workforce.

Keywords: Business Intelligence, Data Warehousing, Generative AI, IS Curriculum, SAP University Alliance, Experiential Learning.

ISSN: 2473-4901
v12 n6678

Bridging the Gap: Transitioning Business Intelligence Curricula from On-Premise Data Warehousing to Cloud Analytics

Wenqi Zhou and Sara Khan

1. INTRODUCTION

Modern Business Intelligence (BI) is undergoing a massive architectural shift, transitioning from static, on-premise data warehousing to dynamic, cloud-based, and AI-augmented analytics (Sasmal, 2024). However, undergraduate BI education faces a profound pedagogical dilemma: how to teach the agility of modern cloud tools while still embedding the foundational concepts of structured data modeling.

Transitioning a curriculum to the cloud presents significant barriers to entry for academic institutions. Cloud-based data warehousing tools are notoriously expensive to license and maintain. They also require a massive amount of front-loaded work from instructors to set up environments, provision accounts, and design datasets that comply with university IT security policies. Furthermore, because modern BI relies on a complex stack of interconnected platforms, there is currently no single textbook that provides a comprehensive, end-to-end curriculum covering both traditional on-premise architecture and modern cloud analytics. Consequently, faculty are often left to piece together disparate tools, risking a fragmented learning experience for students.

This paper addresses these challenges by presenting a replicable pedagogical framework used in a senior-level undergraduate Business Intelligence course. This curriculum relies on experiential learning resources provided by the SAP University Alliance. Notably, this academic year marks the first time the Alliance has offered access to SAP Datasphere, a premier cloud data warehousing solution, for educational use. However, this access is provided on a highly restricted basis. Instructors must provide highly accurate estimates of student account usage, and strict compliance policies dictate that if requested accounts remain unused, the institution risks being denied future access to the platform.

Operating within these stringent constraints, this paper details a multi-phased approach that successfully transitions students through the historical evolution of BI. The curriculum begins with strict, on-premise modeling (SAP BW4/HANA), bridges to cloud data warehousing (SAP Datasphere), and culminates in advanced data mining utilizing SAP Analytics Cloud (SAC). By integrating these different tools into a cohesive, semester-long pipeline, this paper provides a practical roadmap for institutions seeking to modernize their BI offerings without a unified textbook.

2. LITERATURE REVIEW AND PEDAGOGICAL MOTIVATION

2.1 The Evolution to Cloud-Based and AI-Augmented BI

Historically, data warehousing instruction has heavily favored rigid, on-premise enterprise solutions, teaching students to store and manage large volumes of structured data using strict dimensional modeling (Chaudhuri & Dayal, 1997). However, while the fundamental reliance on structured data remains critical for operational reporting, the underlying architecture has undergone a paradigm shift. The modern enterprise is moving away from isolated, on-premise silos toward agile, federated cloud architectures.

Crucially, the integration of Generative AI (GenAI) into BI workflows necessitates this architectural shift. Advanced AI reporting and predictive tools cannot easily operate on isolated legacy systems due to integration bottlenecks and computational latency. Rather, AI-driven solutions require the elasticity and seamless data federation capabilities of cloud-based warehousing (Sasmal, 2024). Even when processing strictly structured data, the underlying infrastructure must live in agile, cloud-hosted environments. This modern architecture allows AI algorithms—such as those embedded in contemporary dashboarding tools—to access, transform, and visualize enterprise data in real-time without the structural friction of legacy systems.

ISSN: 2473-4901
v12 n6678

2.2 The Curriculum Gap and Experiential Learning

Despite this industry mandate, higher education curricula have struggled to keep pace. The IS2020 Competency Model for Undergraduate Programs emphasizes that graduates must understand both data management and emerging technologies (Leidig & Salmela, 2021). Yet, while students receive foundational exposure to legacy data warehousing, the high barrier to entry for designing interconnected cloud environments means they are rarely taught how to build the actual cloud pipelines necessary to support modern analytics. A growing body of information-systems education research documents the value of experiential, vendor-supported instruction for building data-management and analytics competencies: faculty have used free vendor academic portals such as the Teradata University Network to teach database, data-warehousing, and BI topics at low cost (Jukic & Gray, 2008), applied experiential-learning frameworks to enterprise-systems courses (Ruhi, 2016), and integrated SAP University Alliances materials to teach data warehousing with SAP HANA (Wang, 2022).

Because proficiency in cloud architecture, data federation, and hybrid BI platforms is now a core competency that directly impacts marketability (Miller et al., 2024), there is an urgent pedagogical need for modern frameworks. By utilizing vendor-supported academic alliances, educators can leverage experiential learning, allowing students to construct knowledge through hands-on interaction with enterprise-grade cloud software rather than abstract theory. This case study contributes to the literature by demonstrating how faculty can sequence these alliance tools to bridge the gap between legacy on-premise systems and modern cloud intelligence. Table 2 (Section 5.3) maps each course phase to the IS2020 competency areas it primarily develops.

3. COURSE CONTEXT AND TECHNOLOGY STACK

3.1 Course Overview and Objectives

The proposed pedagogical framework was developed and implemented within a senior-level undergraduate Business Intelligence course. The overarching goal of this 15-week course is to provide students with both a conceptual and applied understanding of the analytical tools required to support data-driven decision-making in modern organizations. Rather than isolating data warehousing from data mining, the course treats them as a continuous pipeline. The primary learning objectives require students to understand the foundational architecture of enterprise data warehouses, demonstrate the technical proficiency required to extract, transform, and load (ETL) data across both on-premise and cloud environments, and apply advanced analytical techniques to transform raw data into actionable managerial insights. Table 1 summarizes how the fifteen weeks are distributed across the three phases and their platforms.

Table 1. Fifteen-week course schedule by phase, platform, and focus.

Weeks	Phase	Platform	Primary focus	Student deliverable
1–4	Phase 1: Foundation	SAP BW4/HANA (Eclipse IDE)	OLAP review; dimensional/Star-Schema modeling; InfoObjects, ADSOs, Composite Providers; two-step master-data load; Data Flow, Data Source, Transformations, DTP	Manually built Star Schema on the Global Bike data model
5–9	Phase 2: Cloud Bridge	SAP Datasphere (→ SAC preview)	Logical ER modeling via JSON templates; visual drag-and-drop associations; drill-down hierarchies; on-premise vs. cloud ETL comparison	Federated cloud warehouse with deployed analytic views
10–15	Phase 3: Advanced Analytics	SAP Analytics Cloud; R	Exploratory analysis, visualization, regression, classification; predictive and AI-augmented dashboarding	Predictive dashboard on a real-world scenario (e.g., public-health tracking, fraud detection)

ISSN: 2473-4901
v12 n6678

3.2 A Hybrid Technology Stack

Because no single textbook or software package can adequately cover the breadth of the modern BI pipeline, the course employs a hybrid technology stack that is deliberately sequenced to force students to experience the historical evolution of BI architecture firsthand. The course begins with SAP BW4/HANA, accessed via the Eclipse IDE, to teach traditional, rigid, on-premise dimensional modeling. Here, students manage core architectural elements such as InfoObjects, Advanced DataStore Objects (ADSOs), and Composite Providers. Once this foundation is laid, the curriculum bridges to the cloud using SAP Datasphere. This cloud-native platform introduces students to modern logical design, allowing them to rapidly deploy Entity-Relationship (ER) models using JSON templates and federate data through visual associations. Following this, SAP Analytics Cloud (SAC) is utilized for state-of-the-art data visualization, utilizing the hierarchies built in Datasphere for drill-down dashboarding and further analysis. Statistical analysis in the final phase is performed in R, which students use alongside SAC for regression and classification exercises; R therefore complements, rather than replaces, the SAP analytics layer.

3.3 The Role of Standardized Datasets

To ensure this framework is highly replicable for other instructors, the practical exercises rely heavily on standardized academic datasets provided by the SAP University Alliance, specifically the Global Bike data model. By utilizing these standardized, vendor-supported datasets, faculty can focus their instructional time on teaching the architectural transition between on-premise and cloud environments rather than continuously generating synthetic data that complies with complex university IT requirements.

4. COURSE IMPLEMENTATION: THE PHASED PEDAGOGICAL FRAMEWORK

The course is divided into three distinct phases, mirroring the industry's historical evolution from siloed structured data management to AI-augmented analytics. Figure 1 details the on-premise architecture that anchors Phase 1, and the week-by-week distribution of all three phases is given in Table 1.

4.1 Phase 1: Building the Foundation (Weeks 1–4)

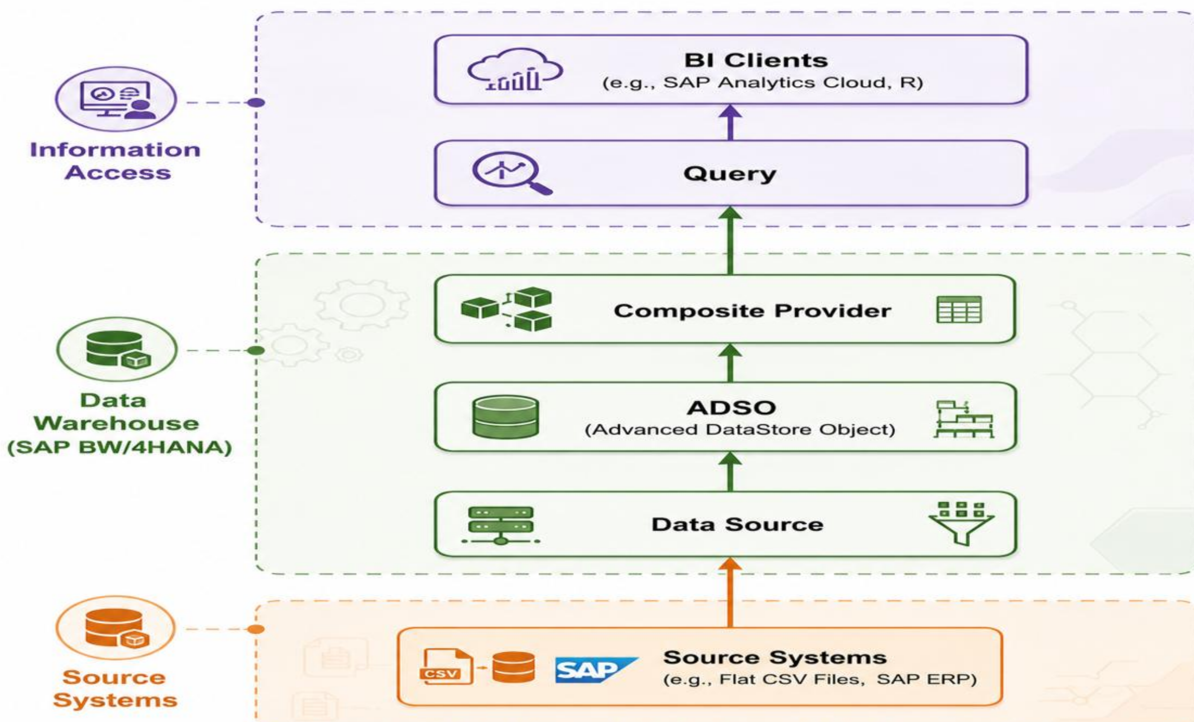

Figure 1. The strict, multi-tiered data flow architecture required in traditional SAP BW4/HANA modeling (Adapted from SAP University Alliances) in the first 4 weeks

 ISSN: 2473-4901
v12 n6678

The first four weeks of the course focus on OLAP review and dimensional data modeling with the pedagogical goal of teaching the strict rules of traditional data architecture. Before students can appreciate the integration capabilities of the cloud, they must understand how enterprise data is structurally managed. Consequently, this phase introduces the fundamental plumbing of Business Intelligence. Students review relational databases and Online Analytical Processing (OLAP) before moving into dimensional data modeling. Figure 1 depicts the strict, multi-tiered data-flow architecture students construct in SAP BW4/HANA during these weeks, tracing structured source data from flat files through the Data Source, Transformation, and Data Transfer Process layers into the dimensional model.

Using the Eclipse IDE connected to SAP BW4/HANA, students manually construct a Star Schema utilizing the structured Global Bike data model. During this phase, instruction focuses heavily on enforcing the strict data typing, extraction, and load rules inherent to on-premise enterprise environments. For example, students learn that loading master data is a rigid, two-step process requiring them to first load attributes (such as Product Category or Division) and subsequently load textual descriptions. To execute this, students must manually define a complete Data Flow, establishing a Data Source, applying Transformations, and configuring a Data Transfer Process (DTP) to execute the extraction from flat CSV files. By mastering these granular mechanics, students gain a profound appreciation for the structural integrity of traditional data warehouses.

4.2 Phase 2: The Cloud Bridge (Weeks 5–9)

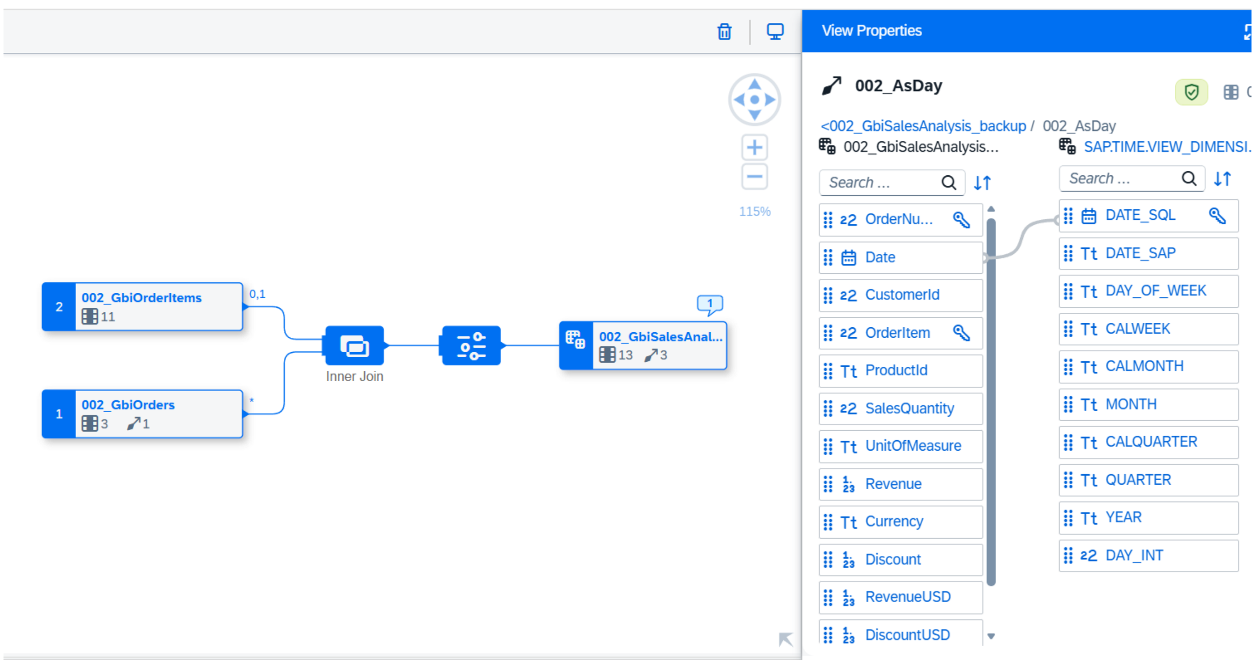

Figure 2. The agile, visual association and mapping interface utilized in SAP Datasphere for modern cloud ETL processes

In the second phase, the curriculum pivots from on-premise architecture to modern cloud computing, focusing heavily on modern Data Extraction, Transformation, and Loading (ETL) processes. Students are introduced to SAP Datasphere, marking their transition into agile data management. Pedagogically, this phase provides a controlled comparison: while the *type* of data remains structured, the *architecture* changes completely.

Instead of the rigid, multi-step DTP configurations required in Eclipse, students in Datasphere learn to rapidly deploy logical ER models using flat JSON file templates. The interface shift is significant; students

 ISSN: 2473-4901
v12 n6678

use visual drag-and-drop mapping to establish joins and associations—such as linking Sales Analysis views directly to Time Dimensions. Furthermore, students build robust Customer and Product dimensions by seamlessly linking attributes and defining drill-down hierarchies. Students directly contrast the rigid implementation requirements they experienced in SAP BW4/HANA with this flexible, federated architecture. Once this cloud warehouse is established and views are deployed, students are introduced to SAP Analytics Cloud (SAC) for initial data exploration, demonstrating how a cloud architecture serves as a necessary bridge to advanced reporting tools. As shown in Figure 2, Datasphere replaces these rigid, code-like configurations with a visual association-and-mapping canvas: the figure illustrates how students drag connections between a Sales Analysis view and its Time and Customer dimensions to build the same logical model they previously assembled step by step in Eclipse.

4.3 Phase 3: Advanced Analytics & Knowledge Discovery (Weeks 10–15)

Having successfully built an on-premise warehouse and transitioned its structured data to the cloud, students spend the final five weeks acting purely as business analysts. The focus shifts entirely to knowledge discovery, utilizing the federated data to generate predictive models and AI-augmented insights. Using SAC, students perform exploratory analysis, data visualization, regression analysis, and classifications on real-world scenarios, such as public health tracking and fraud detection. By culminating the course with predictive analytics, students realize the ultimate value of the complex cloud pipelines they built earlier in the semester, proving that modern AI and machine learning tools are only as effective as the underlying data architectures that support them. In these final weeks students also engage generative-AI capabilities within the analytics layer to accelerate insight generation and to practice prompt engineering against their federated data. In particular, GenAI capacities are explored, such as Smart Insights and Influencing factor analysis, to auto extract data insights and identify root cause. Representative exercises include wine price prediction, fraud detection, and movie recommendation systems.

5. DISCUSSION AND INSTRUCTOR REFLECTIONS

5.1 The Pedagogical Value of the Phased Progression

Reflecting on the implementation of this phased curriculum, the most significant pedagogical success was the deliberate sequencing of legacy systems prior to modern cloud tools. Initially, students often exhibit frustration with the rigid data typing and strict architectural rules required to build a Star Schema in SAP BW4/HANA. However, this struggle is pedagogically necessary. When the curriculum transitions to SAP Datasphere in Phase 2, the contrast is stark; students immediately recognize the value of cloud federation. Because both systems utilize structured data, students experience an "apples-to-apples" architectural comparison. By experiencing the limitations of on-premise silos firsthand, students do not just memorize the definition of "cloud computing"; they genuinely understand the architectural integration problems the cloud was designed to solve (Wixom et al., 2014).

5.2 Managing the Cognitive Load of a Hybrid Stack

The primary challenge of this course design is the high cognitive load placed on students. Because the curriculum eschews a single, unified textbook in favor of a hybrid technology stack, students must rapidly learn and navigate multiple distinct software interfaces (Eclipse, SAP Datasphere, SAC, and R). To mitigate technical fatigue, instructional time must be carefully managed through the lens of experiential learning theory (Kolb, 2014). The instructor must assume the role of a technical facilitator, providing highly structured, step-by-step lab exercises during the initial weeks of each phase. Once the software interface is demystified, the instructor can gradually remove the scaffolding, allowing students to engage in higher-order knowledge discovery and custom prompt engineering in the final weeks. In concrete terms, the early labs supply complete, click-by-click instructions—exact menu paths, field values, and screenshots—so that no student is blocked by an unfamiliar interface. As each phase matures these supports are withdrawn incrementally: step lists give way to goal statements (for example, “federate these two sources and expose a drill-down hierarchy”), and by Phase 3 students frame their own analytical questions, select variables, and interpret model output with minimal procedural guidance, engaging in higher-order knowledge discovery rather than guided replication.

ISSN: 2473-4901
v12 n6678

5.3 Alignment with Industry Demands

Ultimately, this rigorous, multi-platform approach directly addresses the skills gap identified in current IS literature. The modern data analyst rarely works within a single, homogeneous software environment. By forcing students to extract structured data from an on-premise system, federate it in a cloud environment, and push it to an AI-augmented analytical platform (SAC) for analysis, the course mirrors the federated, hybrid reality of modern enterprise IT. Students graduate not merely as operators of a single software package, but as adaptable systems thinkers capable of navigating the complex pipelines required to support AI-driven business intelligence. Table 2 makes this alignment explicit, mapping each phase to the IS2020 competency areas it develops.

Table 2. Alignment of course phases with IS2020 competency areas (Leidig & Salmela, 2021).

Course phase	Platform	IS2020 competency area(s) developed
Phase 1: Foundation	SAP BW4/HANA	Data and Information Management; Systems Analysis & Design (dimensional modeling, ETL rigor)
Phase 2: Cloud Bridge	SAP Datasphere	IT Infrastructure; Data and Information Management (cloud federation, logical modeling)
Phase 3: Advanced Analytics	SAP Analytics Cloud; R	Data/Business Analytics; Emerging Technologies (predictive and AI-augmented analytics)

5.4 Preliminary Evidence and Limitations

To complement these reflections with observational evidence, we examined student performance on a mid-semester quiz of foundational BI concepts—dimensional modeling, ETL, the star schema, the distinction between transactional (ERP) and informational systems, and core SAP data-flow objects—administered in two consecutive offerings of the course (Fall 2024, N = 19; Fall 2025, N = 26). Aggregate mastery was high in both cohorts and improved year over year: the mean score rose from 91.6% to 95.1%, and the proportion of students answering every item correctly rose from 47% to 65%. Restricting the comparison to the nine items whose wording was identical across both offerings, the mean rose from 92.4% to 96.2%. The largest gains appeared on the conceptual-distinction items—separating transactional from informational systems (79% to 96%) and identifying the appropriate system for informational workloads (79% to 88%)—consistent with the framework's aim of building conceptual understanding rather than tool-specific recall. Table 3 reports the item-level results.

Table 3. Item-level performance on the foundational-concepts quiz across two cohorts (identical items only).

Foundational concept assessed	Fall 2024 (N=19)	Fall 2025 (N=26)
Organizing and storing structured data	100%	100%
Identifying a non-database interaction	100%	100%
ERP as a transactional vs. informational system	79%	96%
Appropriate system for informational workloads	79%	88%
ETL process definition	100%	100%
Identifying a non–key figure	95%	96%
Center of a star schema (fact table)	100%	100%
Data acquisition / identification process	89%	96%
Purpose of the Data Transfer Process (DTP)	89%	88%
Overall (nine identical items)	92.4%	96.2%

Anonymous end-of-course student evaluations were consistent with these results. On a five-point agree–disagree scale (5 = strongly agree), students' agreement that “the coursework provided opportunities for me to think critically or in new ways” rose from 4.42 in Fall 2024 to 4.63 in Fall 2025. Although we

 ISSN: 2473-4901
v12 n6678

were concerned that introducing an additional platform and generative-AI activities might overwhelm students, agreement that "I understood what was expected of me in this course" held steady and edged upward, from 4.42 to 4.47, while perceived difficulty remained near the middle of a 0–3 scale (1.92 to 2.00, where 0 = "not challenging enough" and 3 = "too challenging"), indicating that students found the workload appropriately demanding rather than confusing. Taken together, these measures suggest the phased design increased opportunities for critical thinking without eroding clarity of expectations or pushing difficulty to an unmanageable level.

Several limitations bound the transferability of this framework. First, the design is tightly coupled to the SAP University Alliance stack (SAP BW4/HANA, SAP Datasphere, and SAP Analytics Cloud); institutions without Alliance membership, or those standardized on other vendors such as Microsoft, Oracle, Snowflake, or Google Cloud, would need to substitute comparable tools, and the specific interface contrasts that drive the pedagogy may not translate directly. Second, Datasphere access is granted on a restricted, usage-audited basis, which constrains enrollment planning and replicability. Third, the evidence is drawn from a single institution and instructor, limiting generalizability. Fourth, the quantitative evidence reported above rests on modest, unmatched cohorts (N = 19 and N = 26) assessed with a short, recall-oriented instrument on which scores approached a ceiling; it therefore documents attainment and student perception rather than a controlled measure of learning gains, and the two-year improvement cannot be attributed to the curriculum alone.

6. CONCLUSION

As Generative AI and cloud computing reshape the landscape of Business Intelligence, higher education must move beyond teaching data warehousing and data mining as isolated, theoretical concepts. This paper demonstrates that it is feasible to teach the entire BI pipeline, from strict on-premise dimensional modeling to agile, AI-augmented predictive analytics, within a single, semester-long undergraduate course, and it reports the instructors' reflections on why the phased sequence is pedagogically productive; preliminary quiz and student-evaluation evidence (Section 5.4) points to high foundational mastery and increased opportunities for critical thinking, while a controlled evaluation of learning gains remains important future work.

Crucially, by maintaining a consistent focus on structured data throughout the pipeline, this curriculum isolates the architectural shift, allowing students to clearly understand the value of cloud federation. While the barriers to entry for cloud-based education remain high due to software costs and the absence of unified textbooks, leveraging experiential learning resources through vendor partnerships like the SAP University Alliance provides a viable path forward. By adopting a phased pedagogical framework that forces students to experience the architectural evolution of data systems, business schools can graduate highly adaptable professionals equipped to bridge the gap between legacy IT plumbing and the future of artificial intelligence.

7. REFERENCES

Chaudhuri, S., & Dayal, U. (1997). An overview of data warehousing and OLAP technology. *ACM SIGMOD Record, 26*(1), 65–74. https://doi.org/10.1145/248603.248616

Jukic, N., & Gray, P. (2008). Teradata University Network: A no cost web-portal for teaching database, data warehousing, and data-related subjects. Journal of Information Systems Education, 19(4), 395–402.

Kolb, D. A. (2014*). Experiential learning: Experience as the source of learning and development* (2nd ed.). Pearson FT Press.

Leidig, P. M., & Salmela, H. (2021*). IS2020: A competency model for undergraduate programs in information systems.* Joint ACM/AIS IS2020 Task Force. Association for Computing Machinery. https://doi.org/10.1145/3460863

ISSN: 2473-4901
v12 n6678

Miller, L., Dinger, M., Bortz, N., & Mandviwalla, M. (2024). *Information systems job index 2024.* Temple University, Association for Information Systems.

Ruhi, U. (2016). An experiential learning pedagogical framework for enterprise systems education in business schools. The International Journal of Management Education, 14(2), 198–211. https://doi.org/10.1016/j.ijme.2016.04.006

Sasmal, S. (2024). Data warehousing revolution: AI-driven solutions. *International Research Journal of Engineering & Applied Sciences, 12*(1), 1-6. https://doi.org/10.55083/irjeas.2024.v12i01001

Wang, M. (2022). Teaching data warehousing with SAP HANA. Issues in Information Systems, 23(4), 254–264. https://doi.org/10.48009/4_iis_2022_122

Wixom, B. H., Ariyachandra, T., Douglas, D. E., Goul, M., Gupta, B., Iyer, L., Kulkarni, U., Mooney, J. G., Phillips-Wren, G., & Turetken, O. (2014). The current state of business intelligence in academia: The arrival of big data. *Communications of the Association for Information Systems, 34*(1), 1. https://doi.org/10.17705/1CAIS.03401