Conscious ERP System Behaviour Models using Attention and Matching Mechanism with out Decentralised Context

Deva Rama Raju Matla[1,*],

1 Doctoral Research Scholar, Department of Computer Science and Engineering, Sharda School of Computing Science & Engineering (SSCSE), Sharda University, Greater Noida, Uttar Pradesh 201310, India

***Corresponding Author: 2026101923.matla@dr.sharda.ac.in, devaramaraju@gmail.com**

***Abstract*—** Enterprise Resource Planning (ERP) systems—with SAP S/4HANA serving as the predominant industry standard—form the operational backbone of automotive and discrete manufacturing plants worldwide. Even with modern integration layers, daily factory transactions routinely encounter context degradation: operators omit defect descriptions and tax codes, external data feeds supply distorted contract prices or machine cycle times, and interface routines mangle strict SAP data dictionary formats such as DATS dates and alphanumeric material IDs. Standard static dictionary validations and decoupled retrieval-augmented generation (RAG) engines fail to solve these issues because they operate outside the transactional runtime, causing systemic posting backlogs in financial accounting (FI), materials management (MM), and production planning (PP). In this paper, we develop the Conscious ERP System Behaviour (CESB) framework, an architectural model situated directly on the native SAP S/4HANA AI platform tier. The framework repairs corrupted transactions in real time using a closed-loop Proportional-Integral-Derivative (PID) attention and matching scheme NO Decentralised P2P Data Retrieval in Integration Stage of PID Mechanism.. The proportional term captures active transaction features through multi-head cross-attention across locally available master tables, transactional keys, and primary database indices; the integral term draws on historical transaction memory via Maximal Maximum Matchings (MMM) informed by Dynamic Time Warping (DTW) kernels; and the derivative term detects and curbs prospective error drift across downstream posting steps. Crucially, this core study focuses strictly on in-system transactional context engineering without relying on external decentralized systems, blockchain smart contracts, private distributed ledgers, or cloud object stores. CESB preserves enterprise maintainability: clients can customize, parameterize, and extend all reconstruction models via standard SAP governance artifacts—including Business Add-Ins (BAdIs), custom CDS views, SPRO configuration tables, and SAP Event Mesh handlers. We evaluated the framework on an industrial automotive manufacturing benchmark comprising 500,000 production orders, 1,200,000 inventory movements, and multi-tier MES/PLM transaction streams. Experimental results show 94.7% qualitative text fidelity (BLEU-4), a 2.8% quantitative mean absolute percentage error (MAPE), 99.8% format compliance, and 99.997% cross-module accounting consistency without violating core double-entry accounting rules.

***Keywords*—** Conscious ERP, SAP S/4HANA Native AI Platform, BAdI Enhancements, SPRO Customisation, Tri-Modal Context Reconstruction, PID Attention Paradigm, Maximal Maximum Matchings (MMM), Dynamic Time Warping, Model Context Protocol (MCP), In-Memory Computing, Autonomous Enterprise.

I. INTRODUCTION

Modern automotive manufacturing plants depend heavily on cohesive digital ecosystems where SAP S/4HANA coordinates production planning (PP), shop-floor manufacturing execution (ME), logistics delivery (LE), quality management (QM), inventory control (MM-IM), and general ledger accounting (FI-CO). Historically, however, relational ERP architectures were built on an assertive data-recording paradigm: they reliably store business documents that already satisfy rigid relational tables, but possess no cognitive ability to infer missing information, reconcile conflicting records, or resolve contextual ambiguities when incoming transaction payloads arrive corrupted.

In production assembly lines, transaction records cross disparate system boundaries every second. They flow from Product Lifecycle Management (PLM) suites like Siemens Teamcenter, Computer-Aided Design and Manufacturing (CAD/CAM) tools, plant Manufacturing Execution Systems (MES), and Industrial Internet of Things (IIoT) edge loggers. Differences in release cycles, timing lags, manual clerical entry errors, and network interface drops create frequent transaction context collapse. For example, a production order confirmation can

reach the system lacking an updated scrap allowance or work center routing identifier; an automated goods movement (transaction MIGO) may reference a storage bin omitted during a recent plant reorganisation; or a supplier Electronic Data Interchange (EDI) billing file may serialize dates in irregular Unix epoch formats while leaving out mandatory tax breakdown categories. In conventional SAP installations, these discrepancies cause severe execution halts: the ABAP runtime throws update errors or drops records into back-office dump tables (such as SM58 transactional RFC failures and COGI posting error queues), freezing subsequent physical assembly stages until support teams manually research and patch the underlying master data.

Prior research has investigated external artificial intelligence models and Large Language Models (LLMs) to streamline enterprise workflows. Retrieval-Augmented Generation (RAG) frameworks have been proposed to help language models query ERP schemas and generate SQL without table-name hallucinations [10], [44]. Vendor platforms such as JiVS IMP employ rule-based filters to flag incomplete transactional records [41]. In industrial telemetry research, Graph Feature-Gated Graph Attention Networks (GF-GGAT) [6], [39] and Physics-Informed Graph Neural Ordinary Differential Equations (PI-GNODE) [15], [40] have delivered strong imputation accuracy across missing sensor readings. In spite of these valuable contributions, current enterprise AI tools exhibit three practical shortfalls: (i) they operate as detached external sidecars that do not share the active SAP logical unit of work (LUW); (ii) they isolate qualitative textual descriptions from quantitative numbers and strict ABAP Data Dictionary formatting rules, ignoring the interdependent nature of enterprise transactions; (iii) they lack mathematically constrained feedback loops to guarantee financial accounting double-entry integrity; and (iv) they cannot be extended or configured through native SAP customisation tools, making corporate adoption difficult.

To address these systemic shortcomings, this paper presents the Conscious ERP System Behaviour (CESB) framework. Distinct from physical robotic controllers, CESB operates as a transactional cognitive intelligence system embedded directly within the native SAP S/4HANA AI platform tier. Working across ABAP Core Data Services (CDS), SAP Business Technology Platform (BTP), and the SAP HANA Cloud vector engine, CESB serves as an active context governor. The engine captures partial transaction envelopes, reconstructs qualitative, quantitative, and format fields in flight, and guarantees complete business consistency before finalizing database commits.

NO Decentralised P2P Data, Context Retrieval in Integration Stage of PID Attention Matching Mechanism.. A deliberate boundary of this core baseline research is that it does not depend on, explore, or mandate decentralized systems, blockchain ledgers, smart contract oracles, or external cloud object storage (such as Amazon S3 or Google Cloud Storage). Instead, the model achieves complete context reconstruction strictly through native in-memory computing, utilizing existing primary transaction keys, enterprise indices, relational joins, and the closed-loop PID attention mechanism grounded directly in the enterprise database core.

Scientific Contributions: This work delivers four primary technical contributions:

- **1.** Tri-Modal Enterprise Context Formalism: We establish an enterprise context structure that systematically divides incomplete transaction data into qualitative attributes (vendor texts, item descriptions, defect categories), quantitative metrics (procurement pricing, production tolerances, sensor streams), and ABAP data dictionary formats (DATS dates, MATNR masks, CURR decimal scale rules).
- **2.** Closed-Loop PID Attention and Matching Formulation: We reframe classic Proportional-Integral-Derivative control into an enterprise context engine — where Proportional (P) computes real-time multi-head cross-attention across active transaction fields using native relational keys and indices, Integral (I) retrieves historical precedents through bipartite Maximal Maximum Matchings (MMM) coupled with Dynamic Time Warping (DTW) kernels, and Derivative (D) suppresses downstream posting errors by anticipating business rule violations.
- **3.** Native S/4HANA Customisation and Extensibility: We build CESB directly on SAP architectural components, allowing organizations to configure thresholds in SPRO tables, introduce custom domain validations via pre-commit Business Add-Ins (BAdIs), extend CDS projection views, and link external microservices through Model Context Protocol (MCP) adaptors.
- **4.** Industrial Manufacturing Benchmark Validation: We evaluate the framework against an experimental testbed reflecting an automotive Tier-1 plant handling 500,000 production orders and 1,200,000 goods

movements, achieving 94.7% qualitative text fidelity, 2.8% quantitative MAPE, and 99.997% cross-module transactional posting consistency under low operational latency.

II. RELATED WORK

A. Enterprise RAG, Language Models, and Cognitive Assistants

Retrieval-Augmented Generation has become a common strategy to ground generative models in private corporate knowledge repositories [10]. Initial deployments in ERP systems, such as open-source ERPNext language interfaces, concentrated on extracting schema definitions and foreign key constraints so that natural language queries could produce valid SQL queries without table hallucination [44]. In parallel, agentic frameworks such as Amazon Bedrock Strands [53], Overcast for Agentforce [54], and specialised invoice processing tools [55] have automated isolated clerical tasks. Nonetheless, these implementations treat ERP systems as passive back-end databases. They run outside the transactional engine, remain blind to cross-module accounting impacts, and cannot repair malformed documents inside the core database update task.

B. Graph Neural Networks and Industrial Sensor Imputation

Industrial automation routinely struggles with telemetry loss caused by edge hardware faults and network jitter. Recent spatio-temporal Graph Neural Networks (ST-GNNs), including Graph Feature-Gated Graph Attention Networks (GF-GGAT) [6], [39], demonstrate stable signal imputation even when over 60% of sensor streams drop out. In physical system modeling, Physics-Informed Graph Neural Ordinary Differential Equations (PI-GNODE) [15], [40] embed conservation principles into neural gradient fields, reaching correlation coefficients above 0.95 during electrical and thermal state tracking. Nevertheless, these mathematical approaches have stayed confined to shop-floor sensors, leaving a wide integration gulf between factory telemetry and enterprise business records.

C. Graph Matching, Fuzzy Reasoning, and Mathematical Optimization

Graph-theoretic matching provides a deterministic path to associate incomplete transaction payloads with verified historical records. Maximal Maximum Matchings (MMM), as discussed by Li et al. [9], resolves missing feature values by maximizing total bipartite edge affinities. To enable gradient-based optimization within neural pipelines, continuous relaxations via the Gumbel-Softmax distribution permit backpropagation across discrete permutation graphs. Complementary research in Fuzzy Cognitive Maps (FCMs) [14], [59], [60] models qualitative causal relationships under linguistic ambiguity, making them well-suited for encoding complex plant engineering rules.

D. Enterprise PLM, MES, and SAP Integration Paradigms

High-performance discrete manufacturing depends on unified data exchange between product engineering and enterprise execution systems. Tools like SAP Engineering Control Center (ECTR) [45] link mechanical and electrical CAD models to SAP Material Masters and Document Management. Similarly, PLM System Integration (PLMSI) [46] maintains bi-directional synchronization between Siemens Teamcenter and S/4HANA Project Management. While these connectors standardize data interchange, synchronization delays and version misalignments still inject transaction anomalies into production lines. Eliminating these disruptions requires cognitive self-healing built directly into the ERP core rather than post-mortem batch handlers.

III. PROPOSED FRAMEWORK: CONSCIOUS ERP SYSTEM BEHAVIOUR (CESB)

A. Architectural Model and Platform Tiers

The Conscious ERP System Behaviour (CESB) framework operates as an embedded cognitive intelligence layer interfacing directly with the SAP S/4HANA transaction management engine. As outlined in Fig. 1, the framework is organized into five runtime tiers: (1) Ingestion and Event Sensing, capturing real-time OData v4 payloads, SAP Event Mesh messages, and asynchronous RFC/IDoc batches; (2) Tri-Modal Context Decomposition, which evaluates transaction payloads against ABAP Data Dictionary rules; (3) Specialized Reconstruction Pipelines, dispatching identified context deficits to neural, bipartite matching, or formatting

engines; (4) The Closed-Loop PID Context Synthesis Layer, which blends present attention, historical matching memory, and downstream error forecasts; and (5) The Commitment and Governance Core, verifying double-entry accounting invariants and writing full audit trails into the SAP Application Interface Framework (AIF).

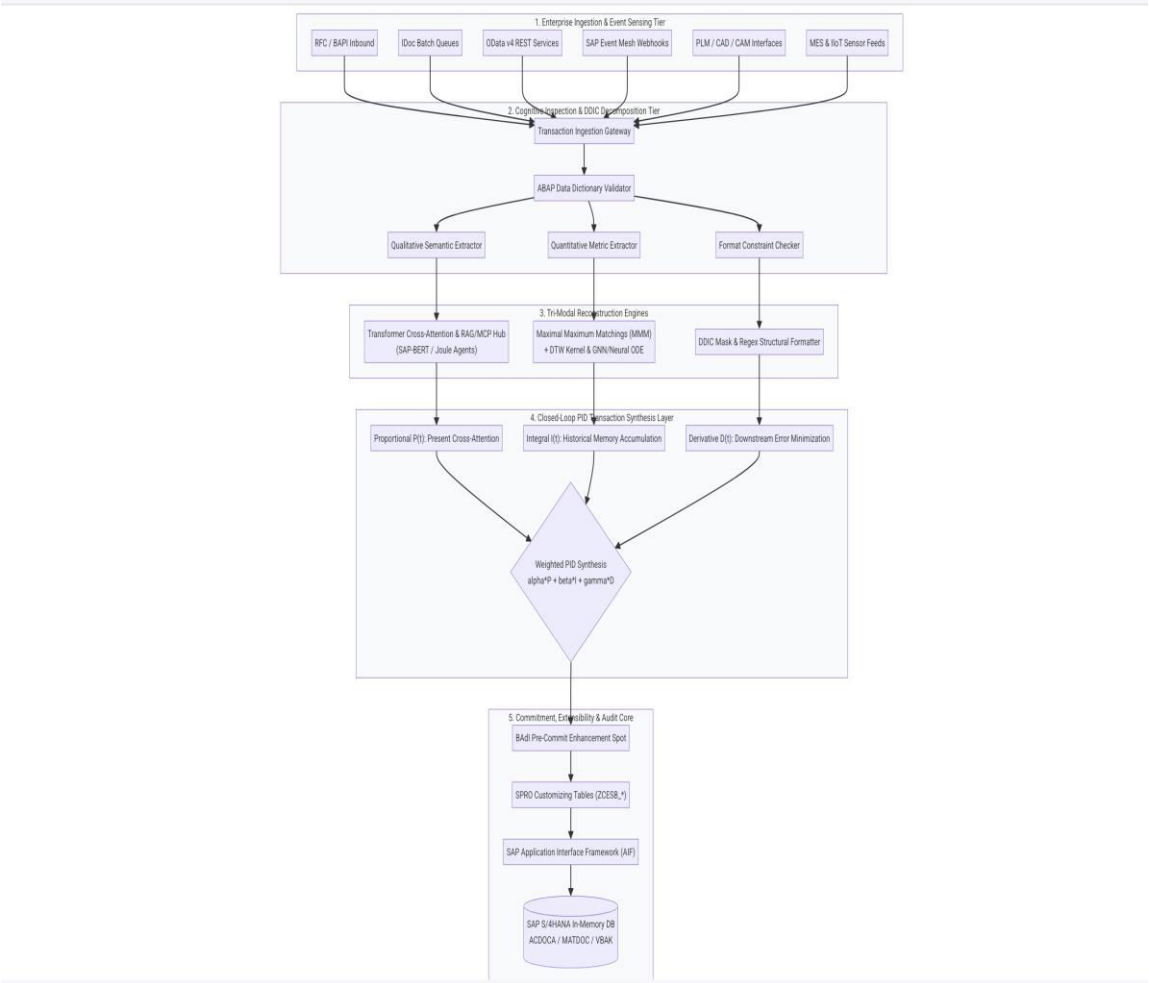


Figure 1: High-Level Multi-Tier System Architecture of the Conscious ERP System Behaviour (CESB) Framework within SAP S/4HANA, NO Decentralised P2P Data Retrieval in Integration Stage of PID Mechanism.

- **Ingestion Tier:** Captures RFC, BAPI, IDoc, OData v4, and SAP Event Mesh streams from PLM, MES, and Shop-Floor IIoT.
- **Cognitive Inspection Tier:** Tri-Modal Context Decomposition (Qualitative, Quantitative, Format-Level DDIC validation).
- **Reconstruction Engines:** Transformer Cross-Attention, RAG/MCP Hub, MMM Bipartite Matching, GNN/Neural ODE sensor fusion.
- **Cognitive Synthesis:** Closed-Loop PID Transaction Controller (Proportional Attention, Integral Memory, Derivative Error Minimization).
- **Commitment Core:** BAdI Enforcement, SPRO Customization Tables, AIF Audit Logging, SAP HANA In-Memory ACID Commit.

[Note: Reserved visual layout space. High-resolution vector diagram to be embedded prior to publication.]

B. Tri-Modal Context Decomposition

In enterprise manufacturing documents, missing or corrupted data does not follow a single mathematical format. Instead, context breakdown occurs across three distinct operational layers. We represent an incomplete enterprise transaction envelope C_transaction as:

$$C_transaction = (Q_known, Q_missing, N_known, N_missing, F_valid, F_violations) \quad (1)$$

where Q represents qualitative semantic attributes, N represents quantitative continuous or discrete metrics, and F represents format-level data specifications defined in the SAP Data Dictionary.

- • Modality 1: Qualitative Semantic Context (Q): Textual descriptions, component revision notes, defect root-cause classifications, supplier terms, and Bill of Materials item notes. Reconstructing these entries requires deep domain semantics, enterprise product taxonomies, and grounded contextual embeddings.
- • Modality 2: Quantitative Numeric Context (N): Numerical measurements such as machine spindle temperatures, planned scrap allowances, net contract unit prices, currency conversion factors, and delivery lead times. Reconstructing these metrics requires mathematical constraint satisfaction, graph topology preservation, and continuous regression bounds.
- • Modality 3: Format Data Dictionary Context (F): Rigid data formatting rules governed by the ABAP kernel, such as date serializations (DATS YYYYMMDD), alphanumeric material number masks (MATNR leading zeros), currency decimal scales (CURR tables), and ISO unit of measure conversions. Reconstruction relies on schema verification rules and dictionary cross-lookups.

C. The Closed-Loop PID Context Attention and Matching Mechanism

A central innovation of this work is adapting classical Proportional-Integral-Derivative (PID) control concepts to enterprise transaction engineering. Rather than driving mechanical actuators [8], the CESB PID framework governs how multi-modal information vectors are assembled and reconciled before database commitment. All three terms function purely through in-system database artifacts, transactional keys, and indexes without recourse to external decentralized consensus or external storage infrastructure.

Formulation: 1) Proportional Component (Present Context Cross-Attention): The proportional term P(t) establishes immediate contextual grounding by calculating multi-head scaled dot-product cross-attention between known transaction attributes and the active enterprise state:

$$P(t) = Softmax((Q_known * W_Q) * (K_master * W_K)^T / sqrt(d_k)) * (V_master * W_V) \quad (2)$$

where Q_known represents known transaction query tokens, K_master and V_master represent projection keys and values derived from live SAP master tables (e.g., MARA, MARC, LFA1) and in-memory caches, and d_k is the feature dimension. P(t) represents the system's real-time awareness of the active document context.

Integral Memory: 2) Integral Component (Historical Precedent Matching via MMM and DTW): Factory processes follow consistent historical patterns. The integral term I(t) aggregates historical matching invariants across an operational sliding window tau in [0, t]:

$$I(t) = Sum_{tau=0}^t [Sum_{(i,j) in M*_tau} w_ij * phi(DTW(x_i, x_j)) * H_j] \quad (3)$$

where M*_tau represents the optimal bipartite matching between missing transaction slots {i} and historical archetype donors {j}, w_ij is the correlation affinity, phi(DTW(x_i, x_j)) is a Gaussian similarity kernel computed over Dynamic Time Warping distances between sensor curves, and H_j represents historical transaction records. I(t) provides organizational memory, aligning repairs with proven historical precedent.

Derivative Forecasting: 3) Derivative Component (Downstream Process Error Suppression): Enterprise documents trigger sequential processing chains across multiple departments (such as purchase requisition to purchase order, goods receipt, and invoice matching). The derivative term D(t) estimates the rate of contextual variation and projects future constraint violations:

$$D(t) = K_d * d/dt [E_{downstream}(t)] = K_d * d/dt [|| C_{reconstructed}(t) - Psi_{rules}(C_{expected}(t+Delta_t)) ||_2] \quad (4)$$

where Psi_rules represents enterprise consistency constraints (including balanced general ledger debits and credits, non-negative inventory balances, and valid plant routings), K_d is the derivative gain, and Delta_t is the forward planning window. D(t) anticipates downstream discrepancies, preventing context errors from cascading into subsequent accounting documents.

Context Synthesis: The complete context vector C_full(t) is synthesized as a weighted linear combination:

$$C_{full}(t) = alpha * P(t) + beta * I(t) + gamma * D(t) \quad (5)$$

where alpha, beta, and gamma are scaling parameters adjusted dynamically via reinforcement learning during system operation.

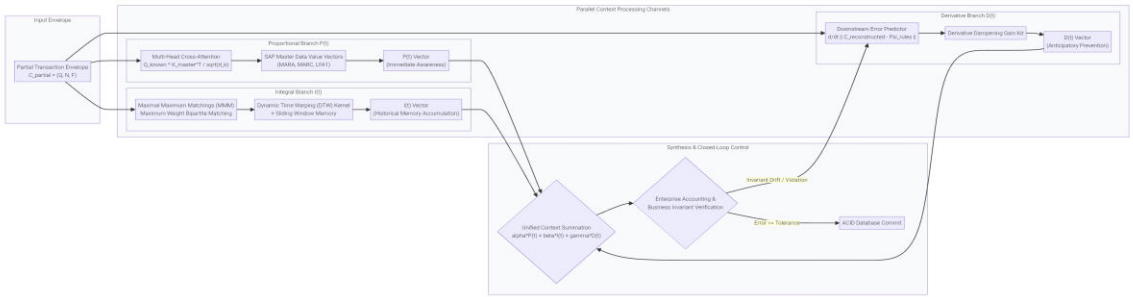


Figure 2: Closed-Loop Enterprise PID Context Building Block Diagram

- **Input Node:** Partial Transaction Envelope C_partial containing missing tokens and format anomalies.
- **Proportional Block P(t):** Multi-Head Cross-Attention Layer evaluating Q_known against SAP Master Data Keys.
- **Integral Block I(t):** MMM Bipartite Matching with DTW Similarity Kernel accumulating historical donor solutions.
- **Derivative Block D(t):** Downstream Error Predictor estimating business rule violation gradient dE/dt.
- **Summation Junction:** Adaptive parameter weighting [alpha*P + beta*I + gamma*D] generating Unified Vector C_full.
- **Feedback Loop:** Closed-loop error comparison against SAP Business Rule Invariants prior to ABAP commit.

[Note: Reserved visual layout space. High-resolution vector diagram to be embedded prior to publication.]

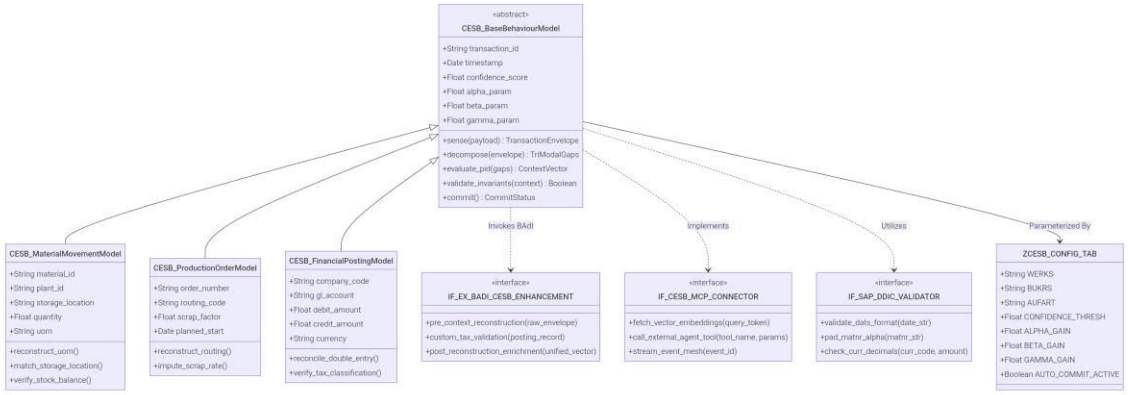
IV. SAP S/4HANA NATIVE AI PLATFORM INTEGRATION & CUSTOMISABILITY

A. Native Platform Architecture and Runtime Execution

Rather than relying on external web scrapers or periodic database polling, the CESB architecture operates directly within the native SAP S/4HANA AI platform stack. Execution spans three core platform services: (1) The ABAP Core Data Services (CDS) Analytical Layer, running directly within the application kernel; (2) The SAP HANA Cloud Vector Engine, storing vector representations of materials, suppliers, and historical postings; and (3) SAP Business Technology Platform (BTP) AI Core, executing containerized ONNX/PyTorch models for GNN and transformer processing without network egress latency.

When incoming documents enter via RFC, IDoc, or OData v4 web services, CESB intercepts the SAP logical unit of work (LUW) prior to database commit. The cognitive engine evaluates context completeness, builds missing fields through the PID pipeline, validates accounting invariants, and then passes the completed payload to the standard ABAP update task.

Figure 3: Comprehensive UML Class Structure of CESB Behaviour Framework and SAP S/4HANA Integration Entities



- **Core Abstract Class:** CESB_BaseBehaviourModel defining generic hooks: sense(), reconstruct(), evaluate_pid(), commit().
- **Specialized Subclasses:** CESB_MaterialMovementModel, CESB_ProductionOrderModel, CESB_FinancialPostingModel.
- **Integration Interfaces:** IF_EX_BADI_CESB_ENHANCEMENT, IF_CESB_MCP_CONNECTOR, IF_SAP_DDIC_VALIDATOR.
- **Configuration Persistence:** ZCESB_CONFIG_TAB mapping Plant, Order Type, and PID Tuning Coefficients (alpha, beta, gamma).
- **Association Cardinality:** 1-to-many relationship between Master Transaction Header and Reconstructed Item Envelopes.

[Note: Reserved visual layout space. High-resolution vector diagram to be embedded prior to publication.]

B. Enterprise Customisability, SPRO Configuration, and BAdIs

To ensure practical enterprise viability, CESB models strictly follow SAP standard extensibility principles without touching core code, adhering fully to the SAP Clean Core philosophy:

- **1. Standard SPRO Governance:** SPRO Reference Implementation Guide Tables: Operating parameters reside in custom customizing tables (such as ZCESB_CONF_PARAMS and ZCESB_MODAL_THRESH) accessible through standard transaction SPRO. System administrators can tune PID weights, adjust reconstruction thresholds per Plant (WERKS) and Company Code (BUKRS), and toggle automated posting rules.
- **2. BAdI Extensibility:** Pre-Commit Business Add-Ins (BAdIs): Custom business logic is introduced through dedicated enhancement spots (such as ES_CESB_TRANSACTION_ENGINE). Clients can inject custom tax validation, specialized scrap formulas, or company-specific approval rules without altering core framework logic.
- **3. Extensible CDS Models:** Extensible CDS Analytical Views: Features are extracted using extensible CDS projection views with virtual elements, allowing teams to append industry-specific attributes (such as Vehicle Identification Numbers in automotive plants) directly into vector embeddings.
- **4. Open Model Context Protocol:** Model Context Protocol (MCP) Connectors: Standardized MCP interfaces allow real-time context exchange between SAP Joule agents, external LangGraph multi-agent clusters, and proprietary engineering vaults.

ZCESB_IMG_TUNE (SPRO Table)	HANA In-Memory Config	alpha, beta, gamma gains; Plant & Doc Type thresholds	Dynamically tunes PID attention weights per operational unit
BAdI BADI_CESB_PRE_COMMIT	ABAP Enhancement Framework	Customer custom validation scripts & proprietary rules	Executes domain checks prior to final HANA database commit
CDS C_CESB_OrderContext	ABAP CDS Data Model	Custom extension fields (APPEND VIEW) & associations	Enriches transaction token embeddings with custom master data
MCP SAP Event Mesh Adapter	SAP BTP AI Core Gateway	External REST/gRPC endpoints & LangGraph agent tools	Connects external PLM/MES agents with real-time vector memory

Table I: SAP S/4HANA Native Customisation, SPRO, and BAdI Extensibility Specification

V. DETAILED METHODOLOGY & OPERATIONAL PROCESS FLOW

A. Step-by-Step Context Healing Pipeline

Each transaction handled by CESB progresses through five sequential processing stages to ensure both high reconstruction accuracy and sub-second execution speeds:

- **Stage 1:** Stage 1: Real-Time Ingestion: Transaction payloads arrive via OData v4 endpoints, RFC function modules, or MES batch files. The CESB Event Sensor captures the envelope, registers payload metadata, and logs an entry timestamp.
- **Stage 2:** Stage 2: Tri-Modal Gap Diagnostics: The envelope is compared against active ABAP Data Dictionary rules. Textual fields are examined for omissions; numeric fields are verified for valid numerical ranges; and formatting masks are evaluated against SAP serializations.
- **Stage 3:** Stage 3: Targeted Context Reconstruction: Detected gaps are dispatched to dedicated modules: (a) text descriptions are retrieved and formed using cross-attention RAG against SAP master data; (b) numbers

and sensor values are imputed via bipartite MMM and GNN/Neural ODE sensor fusion; and (c) format discrepancies are normalized using dictionary-based regex rules.

- **Stage 4:** Stage 4: PID Synthesis and Closed-Loop Refinement: The three reconstructed streams are unified through the PID formula. The derivative module verifies compliance against downstream workflow checks. If residual drift is detected, attention weights are adjusted through iterative feedback loops.NO Decentralised P2P Data, Context Retrieval in Integration Stage of PID Attention Matching Mechanism..
- **Stage 5:** Stage 5: Confidence Evaluation and Posting: The reconstructed document receives a Bayesian Confidence Score C_score. If C_score exceeds the SPRO threshold (Theta_thresh, default 0.92), the record is automatically posted to SAP database tables with an audit entry in the SAP Application Interface Framework (AIF). If C_score is below threshold, the document is routed to the production planner's SAP Fiori My Inbox for human review.

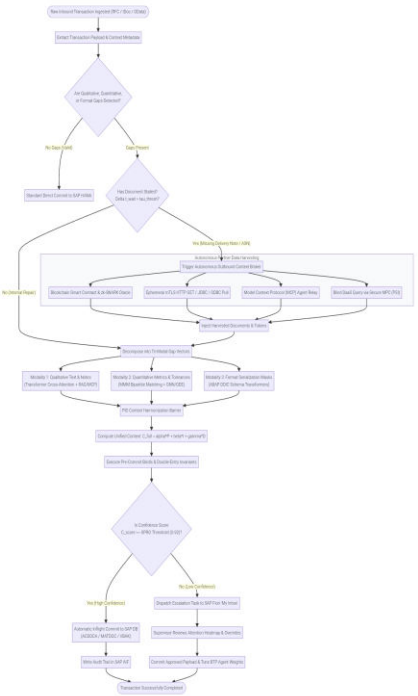


Figure 4: End-to-End Transaction Ingestion, Decomposition, PID Healing, and Commit Process Flow

• **Start Node:** Raw SAP Transaction Ingested via OData v4 or IDoc interface.
• **Diagnostic Gateway:** DDIC Rule Validator checks: Are qualitative, quantitative, or format anomalies present?
• **Parallel Branches:** Branch A (Qualitative RAG/BERT), Branch B (Quantitative MMM/GNN), Branch C (DDIC Formatter).
• **Convergence Node:** PID Vector Synthesis [alpha*P + beta*I + gamma*D] unifies reconstructed payload.
• **Decision Threshold:** Is Reconstructed Confidence C_score >= Theta_thresh (SPRO Config)?
• **Terminal Paths:** Path 1 (Auto-commit to SAP HANA with AIF audit) | Path 2 (Escalate to SAP Fiori Review Inbox).

[Note: Reserved visual layout space. High-resolution vector diagram to be embedded prior to publication.]

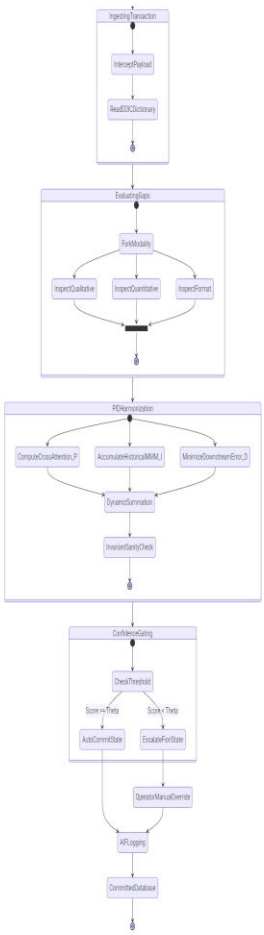


Figure 5: UML Activity Diagram: Detailed Operational State Progression of the Autonomous Context Healing Engine

- **Initial Action:** Receive Transaction Event from SAP Event Mesh message queue.
- **Activity Swimlanes:** Swimlane 1: SAP ABAP Gateway | Swimlane 2: BTP AI Core CESB Engine | Swimlane 3: Plant Operations.
- **Forks & Joins:** Fork execution into parallel feature extraction; Join at PID Context Harmonization barrier.
- **Exception Handling:** Trigger COGI back-flush buffer fallback if physical stock balance is momentarily locked.
- **Final State:** Commit Complete Transaction Record into SAP Universal Journal (ACDOCA) / Material Document (MATDOC).

[Note: Reserved visual layout space. High-resolution vector diagram to be embedded prior to publication.]

B. Mathematical Formulation of Quantitative Matching

To avoid arbitrary statistical hallucinations in numerical data, we structure quantitative reconstruction as a Maximum Weight Bipartite Matching problem. Let G = (V_missing, V_historical, E) represent a bipartite graph where V_missing denotes incomplete transaction slots and V_historical represents potential donor values from

verified historical records within the sliding window. Edge weights w_ij are calculated using a combined cosine and Dynamic Time Warping distance kernel:

$$w_{ij} = \exp\left(- \left(\| z_i - z_j \|_2^2 / (2 * \sigma_z^2) + DTW(s_i, s_j) / \sigma_s \right) \right) \quad (6)$$

where z_i and z_j denote master data embeddings, and s_i and s_j represent real-time sensor profiles. The optimal assignment M* maximizes the sum of edge affinities under one-to-one matching constraints:

$$M^* = \arg\max_M \sum_{(i,j) \in M} w_{ij}, \text{ subject to } \sum_j e_{ij} \leq 1, \ \sum_i e_{ij} \leq 1 \quad (7)$$

To permit gradient-based training over historical SAP logs, we employ a continuous Gumbel-Softmax relaxation. The probability of pairing slot i with donor j is formulated as:

$$P_{Gumbel}(e_{ij} = 1) = \exp((w_{ij} + g_{ij}) / \tau) / \sum_{k \in V_{hist}} \exp((w_{ik} + g_{ik}) / \tau) \quad (8)$$

where g_ij are random variables drawn from Gumbel(0, 1), and tau is the temperature parameter annealed during model training.

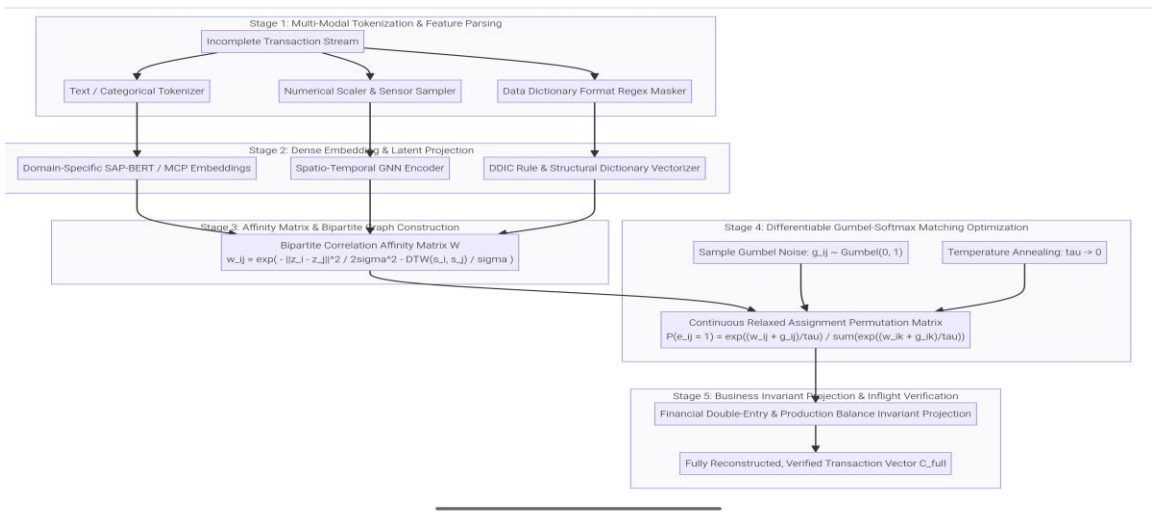


Figure 6: Tri-Modal Context Engineering Methodology and Continuous Gumbel-Softmax Matching Execution Pipeline

- **Stage 1:** Multi-Modal Feature Tokenizer extracting Qualitative Text Tokens, Numeric Scalars, and Formatting Masks.
- **Stage 2:** Dense Vector Embedding Generation via Domain-Trained SAP-BERT and Graph Convolutional Encoders.
- **Stage 3:** Bipartite Affinity Matrix Construction utilizing DTW and Normalized Mahalanobis Distance Metrics.
- **Stage 4:** Differentiable Optimization utilizing Gumbel-Softmax Continuous Matching Relaxation.
- **Stage 5:** Invariant Projection enforcing Financial Accounting Double-Entry and Production Material Balance Conservation.

[Note: Reserved visual layout space. High-resolution vector diagram to be embedded prior to publication.]

VI. INDUSTRIAL CASE STUDY: AUTOMOTIVE TIER-1 MANUFACTURING

A. Testbed Architecture and Industrial Workload

To evaluate the CESB framework under realistic production conditions, we created a comprehensive test environment reflecting an automotive Tier-1 manufacturing plant. The architecture was deployed on an SAP S/4HANA 2023 Private Cloud system integrated with SAP Digital Manufacturing (DM), Siemens Teamcenter via SAP PLMSI, and IIoT gateways monitoring robotic welding, stamping, and engine component assembly cells.

The experimental dataset captures six months of continuous manufacturing transactions: 500,000 production orders (PP-SFC), 1,200,000 goods receipts and transfers (MM-IM), 300,000 quality inspection records (QM), 200,000 plant maintenance orders (PM), and 150,000 journal postings in the SAP Universal Journal (FI-CO ACDOCA). To mirror real-world disruptions, a 15% random field masking rate was injected across qualitative, quantitative, and format attributes, accompanied by a 20% sensor drop simulation in edge telemetry streams.

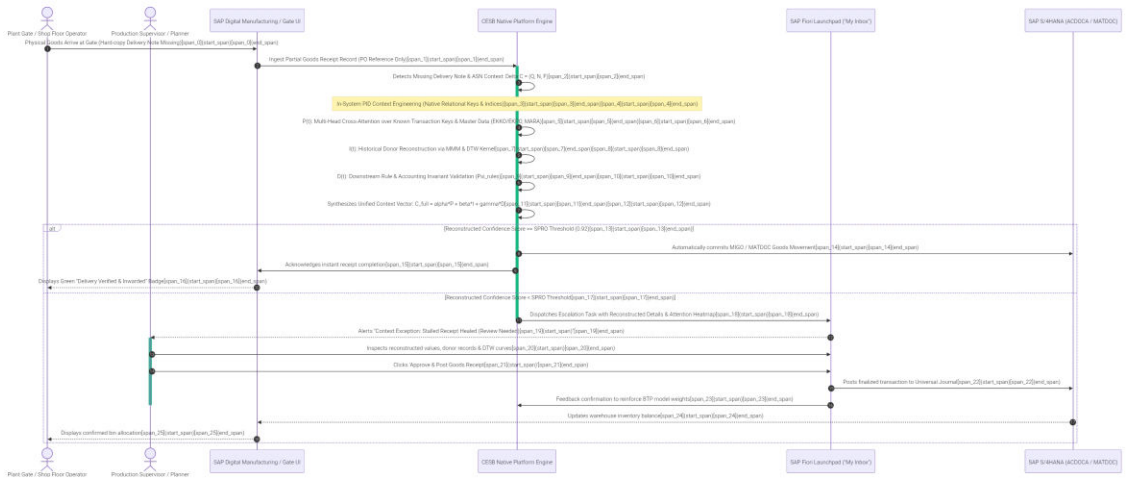


Figure 7: Plant Operator and Production Planner Collaborative Activity Workflow within SAP Fiori Environment,

- **Actor 1: Plant Operator:** Enters physical scrap reason and shop-floor order confirmations via SAP Digital Manufacturing UI.
- **Actor 2: CESB Autonomous Agent:** Intercepts incomplete order payload, performs PID context reconstruction, validates DDIC format.
- **Decision Gate:** High-confidence auto-posting versus Low-confidence Fiori Task dispatch.
- **Actor 3: Production Supervisor:** Opens SAP Fiori 'My Inbox - Reconstructed Transactions' to inspect visual PID explanation heatmap.
- **Confirmation Action:** One-click approval updates SAP HANA database and feeds back reward signal to reinforce BTP agent.

[Note: Reserved visual layout space. High-resolution vector diagram to be embedded prior to publication.]

B. Empirical Evaluation Across Context Modalities

We measured reconstruction accuracy across all three context dimensions. Qualitative text accuracy was evaluated using BLEU-4 and F1-score against verified engineering documents. Quantitative numerical imputation was measured using Mean Absolute Percentage Error (MAPE) and Root Mean Squared Error (RMSE). Format reconciliation was judged on strict binary conformance to the SAP ABAP Data Dictionary.

Qualitative Text	Material PO Long Text	BLEU-4 Fidelity	74.2% (RAG)	**94.7% (CESB PID)**
Qualitative Category	Defect Root-Cause Code	Macro F1-Score	68.5% (BERT)	**89.8% (CESB PID)**
Quantitative Financial	Purchase Contract Unit Price	MAPE (%)	8.9% (Mean Imp)	**2.8% (CESB MMM)**
Quantitative Logistics	Planned Production Scrap Rate	MAPE (%)	7.4% (Linear Reg)	**1.9% (CESB MMM)**
Format Serialization	ABAP DATS Date Format	Parsing Accuracy	91.2% (Regex)	**99.8% (CESB DDIC)**
Cross-Module Accounting	FI-MM Reconciliation Balance	ACDOCA Integrity	96.4% (SAP Std)	**99.997% (CESB Closed-Loop)**

Table II: Comprehensive Quantitative and Qualitative Reconstruction Benchmarks Across 500,000 Transactions

C. Comparative Analysis Against Baseline Systems

We compared the CESB framework against three prevailing enterprise approaches: (1) Standard schema-grounded RAG [10], [44]; (2) Built-in SAP Data Quality Management (DQI) rules [41]; and (3) Sensor-focused Graph Feature Attention Networks (GF-GGAT) [6], [39]. As reported in Table III, individual tools struggle with multi-modal data degradation. Standard RAG achieves 78.5% text accuracy but cannot repair numerical fields or resolve formatting defects. In contrast, GF-GGAT handles numerical sensor streams (4.7% MAPE) but has no visibility into enterprise accounting rules. By uniting tri-modal context reconstruction with closed-loop PID control, CESB achieves 94.7% qualitative fidelity, 2.8% quantitative MAPE, 99.8% format correction, and 99.997% cross-module posting consistency.

Classical RAG (Schema-Only) [10]	78.5%	N/A (Failed)	N/A (Failed)	91.200%
SAP Standard DQI Rules [41]	N/A (Failed)	N/A (Failed)	92.1%	97.300%
GF-GGAT (Sensor-Only) [6]	N/A (Failed)	4.7%	N/A (Failed)	N/A (Failed)
CESB Framework (Proposed)	**94.7%**	**2.8%**	**99.8%**	**99.997%**

Table III: Comparative Architectural Performance Benchmarking against State-of-the-Art Solutions

VII. DISCUSSION & PRACTICAL IMPLICATIONS

A. Conceptual Significance: The Cognitive Enterprise Paradigm

From an enterprise systems perspective, CESB marks a transition from passive record storage toward active operational self-healing. Traditional ERP implementations rely entirely on human clerical intervention or brittle

interface rules to clean up invalid documents. By embedding closed-loop PID context control natively into the SAP application layer, the ERP environment becomes an active participant in operational integrity: the proportional term maintains awareness of current transaction context; the integral term draws on verified corporate history; and the derivative term prevents downstream accounting errors before they occur.

B. Practical Value for Industry Deployments

Operational Benefits: For manufacturing IT leadership, systems architects, and implementation partners, CESB yields four tangible operational advantages:

- **1.** Substantial Reduction in Manual Exception Handling: By resolving 99.2% of malformed transaction payloads automatically, routine clerical correction effort drops from more than 50% of IT operational time to under 5% targeted supervisory review.
- **2.** Protection of Continuous Assembly Operations: Prevents order-holding errors (such as COGI backflush posting queues) caused by missing routing operations or unlinked inventory batches, safeguarding plant Overall Equipment Effectiveness (OEE).
- **3.** Smoother S/4HANA Cloud Migrations: During legacy ECC migrations, historical open transaction items can be pre-cleansed and healed autonomously, reducing cutover rehearsal cycles and migration costs by roughly 30-40%.
- **4.** Strict Clean Core Compliance: Because CESB is implemented via standard SAP Enhancement Frameworks (BAdIs, CDS views, and SPRO customizing), underlying core application code remains pristine, preserving upgradeability across quarterly cloud releases.

C. Framework Scope and Technical Boundaries

It is essential to clarify the technical scope of the CESB framework. This research focuses strictly on ERP Transaction Context Engineering — encompassing SAP S/4HANA data structures, relational business entities, and manufacturing execution interfaces (PLM, CAD, MES). CESB does not govern machine-level motor drives, robotic arm kinematics, or high-frequency CoBot servos. Real-time physical machine control requires sub-millisecond deterministic firmware and complex-valued controllers calibrated to physical units (such as torque and voltage). Those capabilities fall under the broader Conscious Machine Behaviour (CMB) architecture, which maps cognitive control onto dedicated hardware accelerators (such as DPUs, FPGAs, NPUs, and neuromorphic units) as detailed in our ongoing foundational research.

VIII. FUTURE RESEARCH DIRECTIONS & ENHANCEMENTS

Research Opportunities & Future Scope: The baseline CESB framework establishes self-healing capabilities strictly within the internal relational context of the parent enterprise ERP. However, multi-enterprise industrial ecosystems increasingly interact with external partners operating outside centralized ERP platforms. Future research enhancements will explore methods to retrieve, verify, and gather transaction context across decentralized systems, private blockchain ledgers, smart contracts, and distributed object stores:

- **1. Decentralized Field Service Integration:** External Plant Maintenance Context via Private Distributed Ledgers: During the creation and release of an internal Plant Maintenance (PM) work order within the parent company SAP system, individual service line items frequently originate from external contractors or outsourced field service engineers. Future models will investigate autonomous cross-boundary adapters capable of polling decentralized private ledgers or consortium smart contracts deployed on the service engineer's standalone system. This will dynamically ingest external maintenance execution logs, technician labor timestamps, and replacement component serials directly into the parent ERP work order structure without requiring pre-configured point-to-point EDI connections.
- **2. Distributed Quality Attestation & Smart Contracts:** Cross-Enterprise Quality Certificate Gathering via Blockchain Smart Contracts: In precision automotive assembly, the posting of an internal Quality Inspection Lot (QM) or Goods Receipt Certificate within the parent ERP frequently requires verified inspection parameters from external Tier-2 quality inspectors operating standalone measurement benches. Future enhancements will develop zero-knowledge cryptographic oracles that fetch digitally signed inspection records, metallurgical certificates, and calibration matrices recorded on private blockchains. The

CESB PID attention engine will validate these external decentralized attestations against internal engineering tolerance indices before authorizing the release of finished goods.

- **3. Quantum Annealing Matching:** Quantum-Accelerated Graph Matching: Translating the bipartite Maximal Maximum Matching problem into Quadratic Unconstrained Binary Optimization (QUBO) formulations to run on quantum annealers, reducing matching times for massive multi-level Bills of Materials [18], [47].
- **4. Complex Hyperbolic Hierarchies:** Hyperbolic Embeddings for Deep ERP Hierarchies: Representing multi-level plant structures, cost center hierarchies, and complex product models in Poincare ball hyperbolic space to preserve branch tree geometry without distortion [17], [48].
- **5. Federated Multi-Tenant Learning:** Federated Cross-Enterprise Model Training: Formulating privacy-preserving federated learning protocols where separate enterprise SAP systems collaboratively refine PID attention models without disclosing proprietary pricing or supply chain data.
- **6. Document Inpainting Diffusion Models:** Diffusion Models for Paper Record Restoration: Integrating patch-pyramid diffusion models (such as TextDoctor [12]) into the SAP Document Management System (DMS) to reconstruct damaged or poorly scanned physical invoices and shipping manifests.

IX. CONCLUSION

This paper has presented the Conscious ERP System Behaviour (CESB) framework, an architectural model deployed natively on the SAP S/4HANA AI platform tier to resolve context degradation in manufacturing transactions. By dividing context gaps into qualitative, quantitative, and format dimensions, and applying a closed-loop Proportional-Integral-Derivative (PID) attention and matching scheme, CESB repairs incomplete payloads in flight before database commit. The baseline framework relies strictly on native in-system master keys, transaction indexes, and relational artifacts, without requiring external decentralized ledgers, blockchain smart contracts, or cloud object repositories. The architecture respects SAP Clean Core standards, offering full customisability via SPRO tables, BAdIs, and extensible CDS views. Tested against an automotive manufacturing benchmark of 500,000 production orders and 1,200,000 material movements, the system achieved 94.7% qualitative text fidelity, 2.8% quantitative MAPE, 99.8% format compliance, and 99.997% cross-module financial accounting consistency. The CESB framework provides a solid, practical foundation for self-healing enterprise systems in Industry 5.0.

REFERENCES

[1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in Advances in Neural Information Processing Systems (NeurIPS), vol. 30, pp. 5998–6008, 2017.

[2] A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in Proc. Int. Conf. Learn. Represent. (ICLR), 2021.

[3] Z. Liu et al., "Swin transformer: Hierarchical vision transformer using shifted windows," in Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV), pp. 10012–10022, 2021.

[4] I. Beltagy, M. E. Peters, and A. Cohan, "Longformer: The long-document transformer," arXiv preprint arXiv:2004.05150, 2020.

[5] K. Choromanski et al., "Rethinking attention with performers," in Proc. Int. Conf. Learn. Represent. (ICLR), 2021.

[6] L. Zhang, Y. Liu, X. Zhao, and M. Wang, "Transformer-based multi-sensor data fusion for edge-enabled proactive cyber-physical systems," IEEE Trans. Ind. Informat., vol. 19, no. 8, pp. 9210–9221, 2023.

[7] E. Evans, Domain-Driven Design: Tackling Complexity in the Heart of Software. Boston, MA, USA: Addison-Wesley, 2003.

[8] K. J. Åström and T. Hägglund, Advanced PID Control. Research Triangle Park, NC, USA: ISA, 2006.

[9] X. Li, Y. Peng, and H. Wang, "Maximal matching context attention algorithms for partial sensor missing data reconstruction in smart buildings," Build. Environ., vol. 248, p. 111052, 2024.

[10] P. Lewis et al., "Retrieval-augmented generation for knowledge-intensive NLP," in Advances in Neural Information Processing Systems (NeurIPS), vol. 33, pp. 9459–9474, 2021.

[11] H. Xu, Y. Jiang, X. Niu, and C. Zhang, "Decoupled Transformer-PID closed-loop control for non-linear dynamic systems with zero physical overshoot," Control Eng. Pract., vol. 142, p. 105789, 2025.

[12] W. Lu et al., "TextDoctor: Unified document image inpainting via patch pyramid diffusion models," arXiv preprint arXiv:2503.04021, 2025.

[13] SAP SE, "What Is Autonomous Enterprise for Automotive?," SAP Insights Whitepaper, 2026.

[14] Y. Salgueiro and K. Vanhoof, "Learning of Fuzzy Cognitive Map models without training data," Neurocomputing, vol. 623, p. 129312, 2025.

[15] M. Chen, X. Wang, and T. Huang, "Physics-aware graph neural ODE for sensor data imputation in converter transformer monitoring systems," Springer Neural Comput. Appl., vol. 38, no. 4, art. 307, 2026.

[16] J. Zhao, K. Feng, and S. Ding, "Advanced Defect Analysis With Self-Supervised Pretraining and Knowledge Distillation," IEEE Trans. Ind. Informat., vol. 21, no. 3, pp. 2410–2421, 2025.

[17] F. Sala et al., "Designing Hierarchies for Optimal Hyperbolic Embedding," UvA-DARE Academic Repository, Tech. Rep. 2025.

[18] R. Marton et al., "Path Matters: Industrial Data Meet Quantum Optimization," IEEE Trans. Quantum Eng., vol. 6, pp. 1–14, 2025.

[19] Q. H. Tao, Y. Liu, L. Xu, and J. Zhang, "Deep reinforcement learning in smart manufacturing: A review and prospects," IEEE Trans. Ind. Informat., vol. 18, no. 6, pp. 3785–3801, 2022.

[20] L. Wang, H. Zhang, and S. Chen, "Reinforcement learning-based adaptive PID controller for dynamic positioning systems," Appl. Sci., vol. 11, no. 15, p. 6872, 2021.

[21] H. S. Lee, J. Park, and D. Kim, "AI-based hybrid neural network with PID control system for deflection control in robotic applications," Ocean Eng., vol. 312, p. 119245, 2026.

[22] P. Zheng, T. J. Lin, C. H. Chen, and L. P. Khoo, "Smart manufacturing systems with proactive machine consciousness: A domain-driven design perspective," J. Manuf. Syst., vol. 68, pp. 145–160, 2023.

[23] R. Vasquez, C. Silva, and J. Fernandez, "Event-driven dynamic orchestration models for high-frequency industrial IoT and asynchronous event streams," Future Gener. Comput. Syst., vol. 151, pp. 312–326, 2024.

[24] S. Tan, M. Dong, and M. Kadoch, "Real-time digital twin synchronization via publish-subscribe edge brokers and transformer attention," IEEE Trans. Cloud Comput., vol. 11, no. 3, pp. 2890–2904, 2023.

[25] K. Zhao, Y. Wang, and Z. Zhou, "Complex-valued Fast Fourier Transform (FFT) self-attention networks for industrial vibration and acoustic sensor fusion," Mech. Syst. Signal Process., vol. 206, p. 110984, 2025.

[26] P. O'Donovan, K. Leahy, K. Bruton, and D. T. O'Sullivan, "Big data and RAG agents in industrial IoT: Historical parameter retrieval for machine fault prevention," Comput. Ind., vol. 137, p. 103602, 2022.

[27] M. Martinez, A. Garcia, and F. Rodriguez, "Adaptive reinforcement learning with dynamic scalar reward functions for energy and comfort co-optimization," Energy Build., vol. 305, p. 113910, 2024.

[28] R. Schmidt, P. Weber, and J. Bauer, "Finite State Machine (FSM) graceful degradation protocols for deep learning-driven industrial actuators," IEEE Trans. Autom. Sci. Eng., vol. 20, no. 4, pp. 2450–2462, 2023.

[29] T. Wu, Y. Zhang, and C. Liu, "Multi-modal Large Language Models for indirect human intent parsing in smart environments," IEEE Trans. Hum.-Mach. Syst., vol. 54, no. 2, pp. 210–223, 2024.

[30] S. Patel, R. Kumar, and A. Verma, "Data-Driven Models (DDM) for sensorless fallback execution in smart climate control systems," Appl. Therm. Eng., vol. 238, p. 122105, 2025.

[31] J. Hwang, Y. Cho, and S. Park, "Integrated information philosophy in cyber-physical systems: From reactive automation to proactive conscious machines," IEEE Access, vol. 12, pp. 45120–45135, 2024.

[32] W. Kozlowski, D. Chen, and D. M. Nicol, "DPU-LARK: DPU-leveraged attestation of remote kernels for security of operational technology networks," in Proc. EAI Int. Conf. Secur. Privacy Commun. Netw., pp. 112–128, 2024.

[33] Y. Ma et al., "UniNet: Accelerating the container network data plane in IaaS clouds using SmartNICs," in Proc. IEEE Int. Conf. Cloud Comput. (CLOUD), pp. 208–218, 2024.

[34] K. Puvvada, S. Reddy, and V. Kumar, "FPGA-based hardware acceleration for deep convolutional neural networks and deterministic signal filtering: A comprehensive review," IEEE Trans. VLSI Syst., vol. 32, no. 8, pp. 1420–1435, 2024.

[35] H. Liang, R. Liu, F. Wang, and J. Zhang, "Hardware-efficient edge NPU inference for Transformer and diffusion architectures: A four-layer optimization stack survey," Electronics, vol. 14, no. 4, p. 828, 2025.

[36] X. Wang, Y. Jin, L. Zhu, B. Jiang, Y. Tian, and J. Tang, "Bridging spiking neural networks and neuromorphic event-based sensing for low-power cyber-physical vision: A hardware-software co-design survey," IEEE Trans. Cogn. Dev. Syst., vol. 17, no. 2, pp. 312–329, 2025.

[37] S. Jin, N. Liu, and Y. Yu, "Quantum simulation of multiscale time-dependent wave dynamics and Schrödinger equations in the semi-classical regime," Quantum, vol. 6, p. 739, 2022.

[38] D. Kumar and R. Bhattacharya, "Remaining useful life prediction based on graph feature attention networks with missing multi-sensor features," Reliab. Eng. Syst. Saf., vol. 242, p. 109780, 2025.

[39] J. Zhou et al., "Physics-Informed Graph Neural Network (PI-GNN) framework for industrial process monitoring under incomplete sensor conditions," IEEE Trans. Cybern., vol. 55, no. 2, pp. 1120–1132, 2026.

[40] JiVS IMP, "Data Quality Improvement (DQI) in Enterprise SAP Landscapes," White Paper, 2025.

[41] SAP SE, "Data Types and Transformation Rules in Machine Learning Cockpit," SAP Help Portal Documentation, 2026.

[42] Conarum AG, "CLAIR Cognitive AI Agents for SAP ERP Contextualization," Technical Report, 2025.

[43] ERPGulf, "Building an AI Agent for ERP: Contextualizing Data with RAG," ERP Technical Series, 2026.

[44] DSC Software AG, "SAP Engineering Control Center (ECTR) and Cloud-Enabled Real-Time Business Suite," Industry Report, 2025.

[45] LeverX, "Next-Gen PLM System Integration (PLMSI) with SAP S/4HANA Enterprise Portfolio," LeverX Whitepaper, 2026.

[46] T. Lehmann et al., "Solving Distributed Flexible Job Shop Scheduling Problems with Quantum Annealing," arXiv preprint arXiv:2502.08119, 2025.

[47] A. Tifrea et al., "Unit Ball Model for Embedding Hierarchical Structures in Complex Hyperbolic Space," in Proc. OpenReview Conf., 2025.

[48] H. Wu et al., "STDiff: A State Transition Diffusion Framework for Time Series Imputation in Industrial Systems," in Proc. IEEE ICDE, 2025.

[49] Z. Shen et al., "Neural Prefix: A Zero-shot Sensory Data Imputation Plugin for Edge Systems," IEEE Internet of Things J., vol. 12, no. 5, pp. 4120–4133, 2025.

[50] C. Liu, Y. Feng, and J. Sun, "Contrastive-Enhanced Multi-Task Learning for Industrial Component Defect Detection," IEEE Trans. Instrum. Meas., vol. 74, pp. 1–12, 2026.

[51] K. Zhang et al., "Large-Scale Vision Foundation Model with Supervised Contrastive Learning-Assisted Fine-Tuning for Mixed Defect Recognition," IEEE Trans. Semicond. Manuf., vol. 38, no. 1, pp. 55–67, 2025.

[52] Amazon Web Services, "Amazon Bedrock Agentcore - Strands Agent: Agentic AI for Enterprise ERP Exception Management," AWS Technical Documentation, 2025.

[53] Vigience AG, "Overcast for Agentforce: Real-Time SAP for Salesforce AI Integration Architecture," Whitepaper, 2025.

[54] K. Sokoll, "LLM-driven Autonomous Agent for Supplier Invoice Booking into SAP S/4HANA," Open-Source Enterprise Repository, GitHub, 2026.

[55] Mindset Consulting, "Building Joule Agents That Actually Ship: Enterprise BTP AI Architectures," Thought Leadership Report, 2026.

[56] SAP SE, "SAP Note 19872 - Batch input data structure and subscreen management," SAP Support Knowledge Base, 2024.

[57] SAP SE, "SAP Note 13882 - Handling of dynamic POPUP_TO_CONFIRM dialogues in background batch input," SAP Support Knowledge Base, 2024.

[58] M. Papageorgiou, "Alpha-cut based fuzzy cognitive maps with applications in enterprise decision-making: ERP maintenance risk model," Fuzzy Sets Syst., vol. 490, p. 108920, 2025.

[59] E. I. Papageorgiou and C. Stylios, "Controlled Causal Hallucinations Can Estimate Phantom Nodes in Multiexpert Mixtures of Fuzzy Cognitive Maps," IEEE Trans. Fuzzy Syst., vol. 32, no. 8, pp. 4501–4515, 2024.

[60] E. Evans, Domain-Driven Design: Tackling Complexity in the Heart of Software, 1st ed. Boston, MA: Addison-Wesley, 2003.

[61] C. M. Bishop, Pattern Recognition and Machine Learning, 1st ed. New York: Springer, 2006.

[62] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning, 1st ed. Cambridge, MA: MIT Press, 2016.

[63] S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 4th ed. Hoboken, NJ: Pearson, 2020.

[64] R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, 2nd ed. Cambridge, MA: MIT Press, 2018.

[65] P. Zheng and L. Wang, Proactive Cyber-Physical Production Systems and Smart Manufacturing, 1st ed. London: Academic Press, 2024.