Frontiers in Emerging Engineering & Technologies (Open Access)

Volume 03, Issue 02, February 2026,
DOI: - https://doi.org/10.64917/feet/Volume03Issue02-01 PageNo.01-24

# Conceptual Framework of LLM-Based Copilots in EPC Firms for Automated P&ID Generation to Reduce Design Time and Increase Standardization

**Rohit Shinde**
Lamar University, University of Texas at Austin, USA
**E-mail:** rohit.shinde2290@gmail.com

**Snehal Bhosale**
Lamar University,
**E-mail:** snehal1991bhosale@gmail.com

**RECEIVED** - 01-03-2026, **RECEIVED REVISED VERSION** - 01-15-2026, **ACCEPTED**- 02-20-2026, **PUBLISHED**- 02-26-2026

**Abstract**

Traditional P&ID workflows in EPC firms are labor-intensive and error-prone, causing design inconsistencies that impact safety and efficiency. This analytical study proposes a conceptual framework for integrating Large Language Model (LLM)-based copilots to automate P&ID generation, enhance design efficiency, and ensure adherence to industry standards. The methodology details a rigorous conceptual training regimen and a multi-layered system architecture, including a Chatbot Interface, Planning Agent, Graph-RAG Knowledge Base, Execution Module, and Visualization. Modeled outcomes, derived from comparative analysis using industry benchmarks, anticipate significant efficiency gains, including a 30.5x overall productivity improvement factor, 96.7% average time savings for individual elements, and a 3.1x project acceleration. Projected error reductions are substantial, with critical error rates for missing safety devices estimated to decrease from 9.2% to 0.9% in a simulated environment. These findings suggest LLM-based copilots can transform manual P&ID drafting, mitigating project inefficiencies. Successful AI adoption, however, hinges on resolving challenges related to data security, integration costs, and legal liabilities for safety-critical designs, moving beyond conceptual feasibility towards industrial-grade reliability and robust human oversight.

**Keywords:** Large Language Models (LLMs), Piping and Instrumentation Diagrams (P&ID), P&ID Automation, Graph-RAG, Retrieval-Augmented Generation, EPC Industry, Process Design Engineering, Digital Transformation, Design Automation, DEXPI XML, Human-in-the-Loop (HITL), Process Safety, Engineering Design Efficiency, Knowledge Graphs, Generative AI

## 1. Introduction:

Imagine asking an AI: "Design a pump loop with PSV and control valves with instruments and control signals"—and receiving a ready-to-edit P&ID. This level of automation is rapidly transforming industries such as software development and content creation; however, the chemical engineering design space, particularly within Engineering, Procurement, and Construction (EPC) firms, remains largely untouched by such innovations.

Despite its foundational role in global infrastructure, the chemical process industry (CPI) has notably lagged in adopting digital transformation, particularly in process design. Piping and Instrumentation Diagrams (P&IDs), which are definitive schematics for process plants, are predominantly created through laborious manual workflows. These traditional methods are susceptible to inefficiencies, inconsistencies across design teams, and significant rework owing to their reliance on individual expertise and limited standardization. Even with existing digital tools, engineers spend redundant hours recreating standard configurations, indicating a clear need for more intelligent assistance in the early stage design.

The EPC industry faces critical challenges in P&ID Design, which significantly impact project timelines, costs, and safety. These include (1) Laborious Repetitive Work, where 65–70% of engineering time is spent on manual drafting and compliance checks (Junhyung Byun et al. 2025) (2) Design Inconsistencies, leading to 25–30% of P&IDs requiring rework due to symbol/tag mismatches (Schlegl, T et al. 2022). (3) High Error Rates, with safety-critical errors (e.g., missing PSVs) occurring in 8–12% of manual designs (Dzhusupova, Z et al., 2024).

Although recent digital advancements have primarily focused on data management and post-design reporting, the application of artificial intelligence in generative process design remains limited. The current literature identifies a gap in AI-driven copilots that can interactively generate designs while integrating industry standards (e.g., ISA 5.1 and ISO 10628), thermodynamic principles, and rule-based assemblies. This leads to the central research question: How can Large Language Model (LLM)-based copilots be conceptually integrated into EPC workflows to automate P&ID generation, enhance design efficiency, and ensure adherence to industry standards?

This analytical study proposes a conceptual framework to address this gap and explores how LLMs can be theoretically trained on standard design principles and commonly used P&ID assemblies to assist engineers. It outlines a conceptual framework for integrating an AI-powered copilot interface into existing design software, such as SmartPlant P&ID, enabling engineers to interactively prompt the generation of typical equipment configurations using natural language input. By doing so, this study advocates for a new paradigm in engineering design, where LLM copilots reduce repetitive workloads, improve standardization, and accelerate the early-stage design process in EPC firms, ultimately aiming to provide more accurate designs and avoid industrial hazards caused by design issues.

The remainder of this paper is structured as follows: Section 2 provides a comprehensive Literature Review on AI in engineering design, existing P&ID automation tools, design patterns, AI copilots in enterprise software, and identified research gaps. Section 3 details the conceptual methodology, including the proposed copilot model training, system architecture, example use case, and evaluation metrics. Section 4 presents the anticipated results, focusing on conceptual time savings, error reduction, and design consistency. Section 5 discusses the implications, challenges, and future potential of LLM-based copilot systems in EPC. Finally, Section 6 concludes the paper by summarizing the key findings and outlining the future research directions.

## 2. Literature Review

The digital transformation within the Engineering, Procurement, and Construction (EPC) sector has been a decades-long endeavor, consistently driven by the imperative to overcome persistent inefficiencies inherent in large-scale project execution. This evolution signifies a fundamental shift, moving beyond technology as a mere tool for documentation to its role as a cognitive partner for optimization and integrated project delivery. This foundational shift sets the stage for the advanced, data-driven automation now offered by artificial intelligence.

### *The Evolution of Engineering Design Automation towards AI-Augmented Systems*

The integration of Artificial Intelligence (AI), particularly Large Language Models (LLMs), into engineering design represents a paradigm shift, moving beyond mere task automation to a cognitive partnership between engineers and machines. LLM-based systems have demonstrated significant performance improvements in various industrial applications, with studies indicating substantial error reduction and considerable enhancements in design time compared to conventional approaches for similar tasks (Werheid et al., 2024). These AI technologies leverage machine learning models trained on vast datasets to identify patterns, predict optimal solutions, and propose strategies for new engineering challenges. Generative design algorithms further extend these capabilities by exploring extensive solution spaces that human designers would find impractical to consider manually.

LLMs are being increasingly explored for application across diverse engineering disciplines. For instance, frameworks like "Vibe Engineering" propose utilizing LLMs to automate discipline-specific tasks, such as generating piping layouts or performing electrical load calculations, based on natural language prompts. Such integrated approaches are projected to yield dramatic efficiency gains, including a 30-50% reduction in design cycles and a 70% reduction in costly rework by harmonizing outputs and preempting conflicts across disciplines (Ghosh, 2025). This broad applicability in industrial automation highlights the direct parallels to the EPC design processes.

### *The Emergence of LLM-Powered Copilots and Agentic Workflows for* De Novo *P&ID Generation*

Building upon the broader trends in AI-augmented design, cutting-edge research has shifted towards assisting in the creation of new designs, specifically the de novo generation of Piping and Instrumentation Diagrams (P&IDs). This area represents a significant advancement over earlier literature, which predominantly focused on digitizing existing diagrams. This earlier foundational work relied heavily on computer vision techniques, such as the use of Convolutional Neural Networks (CNNs) for pixel-based layer segmentation and object recognition in complex engineering drawings, a necessary precursor to structured data extraction (Moreno-Garcia et al., n.d.). Pioneering work, such as the ACPID Copilot by (Gowaikar et al., 2024), directly addresses the de novo creation of P&IDs from natural language descriptions. Their methodology employs a multi-step agentic workflow that breaks down complex P&ID generation into manageable, verifiable stages, culminating in an intermediate textual representation (DEXPI XML) before graphical visualization (Gowaikar et al., 2024). This agentic approach, which mirrors human engineering processes for auditable steps and error detection, offers a robust response to the limitations inherent in monolithic LLM generation.

In a comparative context, (Werheid et al., 2024) developed a "factual-driven copilot" for manufacturing equipment selection, which also utilized a multi-agent architecture with Retrieval-Augmented Generation (RAG). While the work by (Werheid et al., 2024) focused on component selection for manufacturing and demonstrated success in providing actionable recommendations, (Gowaikar et al., 2024) work directly tackles generative design for P&IDs, showcasing the feasibility of creating new schematics from natural language. Both studies, however, collectively underscore the efficacy of agentic workflows in enhancing reliability and traceability for complex engineering tasks, particularly in environments demanding adherence to industry standards like ISA 5.1 (Instrumentation Symbols and Identification) and ISO 10628 (Flow diagrams for process plants – General rules).

### *Critical Challenges and Constraints for LLM Deployment in Safety-Critical Engineering*

Despite the considerable promise of LLM-powered copilots, this area of research also rigorously assesses the significant barriers that currently impede their widespread and reliable deployment in industrial engineering contexts. The most fundamental issue is the inherent unreliability and phenomenon of "hallucination" in LLMs. As probabilistic text predictors, LLMs lack a true understanding of facts or physical laws, which contrasts sharply with the deterministic and safety-critical requirements of engineering design, where an incorrect specification could lead to catastrophic failure. The rigidity of established industry standards, such as ISA 5.1 and ISO 10628, poses unique and complex constraints on LLM outputs. Current solutions to ensure compliance often involve post-generation checks utilizing rule-based systems that validate P&ID schematics structured as formal graphs (Schulze Balhorn et al., 2025). While such rule-based autocorrections are valuable, they are inherently limited by the explicit rules they are programmed with and struggle with novel design interpretations. The models' "knowledge" of these codes is merely a statistical association of words, not an executable rule set, making direct, unconstrained generation fundamentally untenable in safety-critical applications.

Furthermore, significant risks exist concerning data security, privacy, and intellectual property (IP). These include vulnerabilities such as prompt injection, sensitive information disclosure, and data/model poisoning, all of which highlight the inherent opacity and security challenges of LLMs. These compounded challenges, particularly the architectural mismatch between LLMs' probabilistic nature and engineering's deterministic constraint satisfaction, underscore the urgent need for novel system designs that extend beyond pure LLM applications to ensure practical and safe integration.

### *Mitigation Strategies, Human-in-the-Loop (HITL) Validation, and the Evolving Role of the Engineer*

In direct response to the critical challenges outlined in the preceding theme, the literature details a suite of mitigation strategies and governance frameworks.

Technical grounding methods, such as Retrieval-Augmented Generation (RAG) and Graph-RAG, are emphasized for combating hallucination by grounding LLM responses in trusted, external knowledge bases like structured knowledge graphs (Han et al. 2024). Complementary approaches include fine-tuning on domain-specific datasets and advanced prompt engineering (e.g., few-shot, chain-of-thought prompting), which aim to enhance accuracy and align model behavior with engineering requirements (Almohaimeed, S et al. 2024).

Crucially, the literature universally emphasizes the indispensable role of robust human oversight through Human-in-the-Loop (HITL) or Human-on-the-Loop systems (Kathiresan, G. 2025). For safety-critical applications such as P&ID generation, HITL is considered essential, not merely as a temporary training scaffold, but as a permanent, legally necessary architectural component to maintain accountability, as AI models cannot be held liable for design failures. This integration fundamentally shifts the engineer's role from content generator to validator, curator, and steward of AI-generated proposals. While LLMs offer significant short-term productivity gains, this theme also explores the "paradox of productivity," where speed can come with new cognitive burdens of verification and the risk of skill atrophy due to over-reliance, highlighting the need for holistic evaluation frameworks that account for both efficiency and long-term human capabilities.

**Comparison to Existing Research**

While existing literature has explored the application of AI in engineering, particularly for design automation, few studies directly address the specific challenge of integrating Large Language Models (LLMs) for automated P&ID generation in the EPC context. For instance, focused on optimizing process flow diagrams using rule-based systems, but their approach lacks the generative and reasoning capabilities inherent in an LLM-based system. Similarly, other research has applied computer vision to interpret existing P&ID drawings , but they do not provide a framework for creating new diagrams from unstructured textual data, as your framework does.

**Gaps in Current Literature**

Despite the demonstrated potential and ongoing advancements, a systematic review of the literature reveals several critical gaps that persist between laboratory demonstrations and the development of robust, reliable, and safe LLM-based systems suitable for industrial deployment in the EPC sector.

**Gaps in Current Literature:**

- Limited EPC-Specific Research: There is a notable scarcity of research specifically focused on LLM-based copilots within the EPC industry context. Most existing studies address general engineering applications or specific technical domains, without fully considering the unique challenges and requirements of EPC project environments (Gowaikar et al., 2024).
- Validation and Reliability Studies: Current literature lacks comprehensive validation studies that robustly demonstrate the reliability and accuracy of LLM-based systems in real-world EPC applications. Much research remains at the proof-of-concept stage, without extensive industrial validation or long-term performance assessment (Werheid et al., 2024).
- Integration with Existing Systems: Limited research adequately addresses the complex integration challenges of implementing LLM-based copilots within existing, often proprietary, EPC software ecosystems and established workflow processes. The compatibility with legacy systems requires further investigation (Liu, X. et al., 2025).
- Economic Impact Assessment: There is insufficient research on the tangible economic implications and return on investment (ROI) of implementing LLM-based copilots in EPC firms. While technical improvements are often demonstrated, comprehensive cost-benefit analyses are lacking (Ya-Ting Chuang et al., 2025).

## 3. Methodology:

This study presents a conceptual methodology for developing an LLM-based AI copilot aimed at automating Process and Instrumentation Diagram (P&ID) generation and enhancing design efficiency in Engineering, Procurement, and Construction (EPC) environments. The approach addresses core challenges in traditional chemical plant design workflows, such as labor-intensive tasks, design inconsistencies, and high susceptibility to errors.

The proposed solution leverages Large Language Models (LLMs) fine-tuned on domain-specific engineering data, and integrates them with a graph-enhanced retrieval-augmented knowledge base to create an intelligent copilot system for process engineers. This system transforms natural language design commands into accurate, standards-compliant P&ID outputs by orchestrating the interaction between user input, AI-driven command interpretation, knowledge-grounded validation, and CAD-based visualization tools.

### 3.1 Conceptual Workflow

The complete operational flow of the system is visualized in Figure 1, which outlines the stepwise transformation from the user prompt to the final P&ID delivery. This workflow is designed to enforce traceability and compliance, highlighting the critical validation and human oversight steps necessary to ensure reliability in a safety-critical domain.

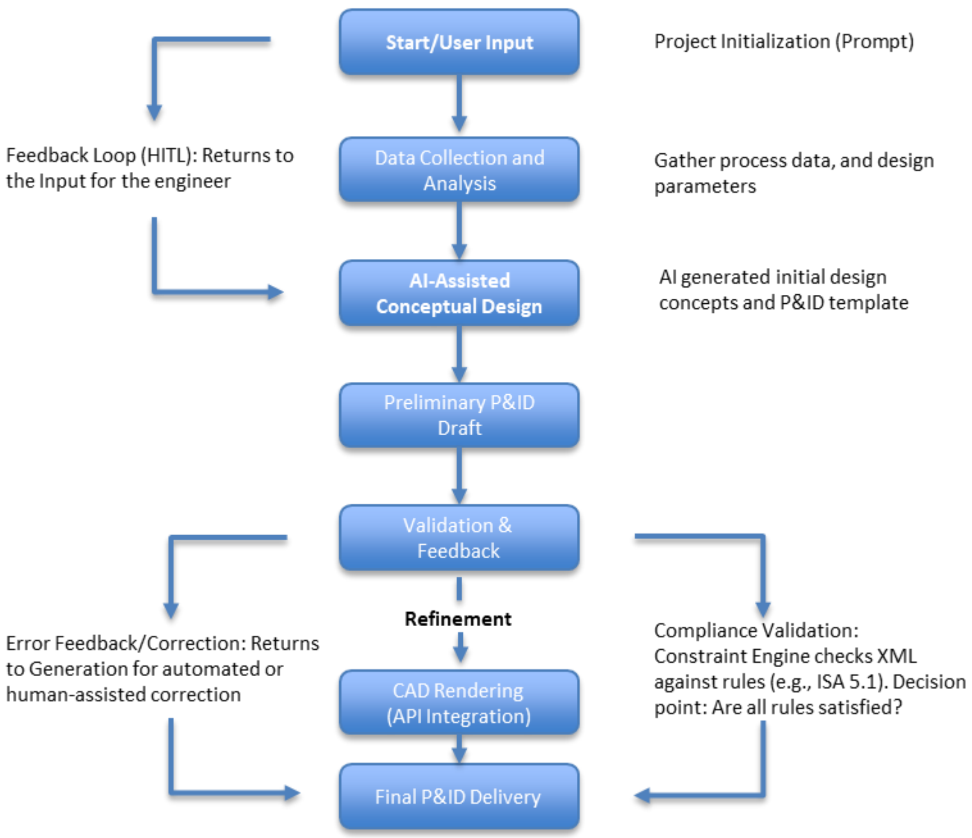


**Figure 1:** Conceptual Copilot Workflow for P&ID Generation

The process begins with the Start/User Input phase, where the engineer provides Natural Language commands. The system proceeds through Data Collection and Analysis and AI-Assisted Conceptual Design, which generates a Preliminary P&ID Draft. Crucially, the draft must pass through a Validation & Feedback gate, where the Constraint Engine enforces rule-based compliance checks. Upon successful validation, the design proceeds to CAD Rendering (API Integration) and finally to the Final P&ID Delivery, which serves as the ultimate Human-in-the-Loop (HITL) gate for engineer sign-off.

### 3.2 System Architecture (Layers and Components)

The system is architected into three distinct layers, as shown in Figure 2, to facilitate modularity, traceability, and seamless integration with existing EPC design tools:

- **User Interface Layer**: Captures natural language input from process engineers via a chatbot interface integrated into the design environment.
- **AI Copilot Core Layer**: Comprises the core intelligence modules—including a command interpreter/planning agent, domain-specific knowledge base (Graph-RAG), and execution/generation module responsible for

generating structured P&ID representations DSL (Domain-Specific Language) or DEXPI XML (Digital Exchange of Piping and Instrumentation).

- **CAD Integration Layer**: Interfaces with industry-standard P&ID software (e.g., SmartPlant or AutoCAD) to render validated graphical outputs and support agentic learning through feedback mechanisms.

This architecture emphasizes modularity, traceability, and engineering standard compliance, while enabling a closed-loop feedback system for continuous refinement. The complete workflow, including feedback integration and human-in-the-loop learning, is visualized in Figure 2, which outlines the stepwise transformation from user prompt to P&ID generation.

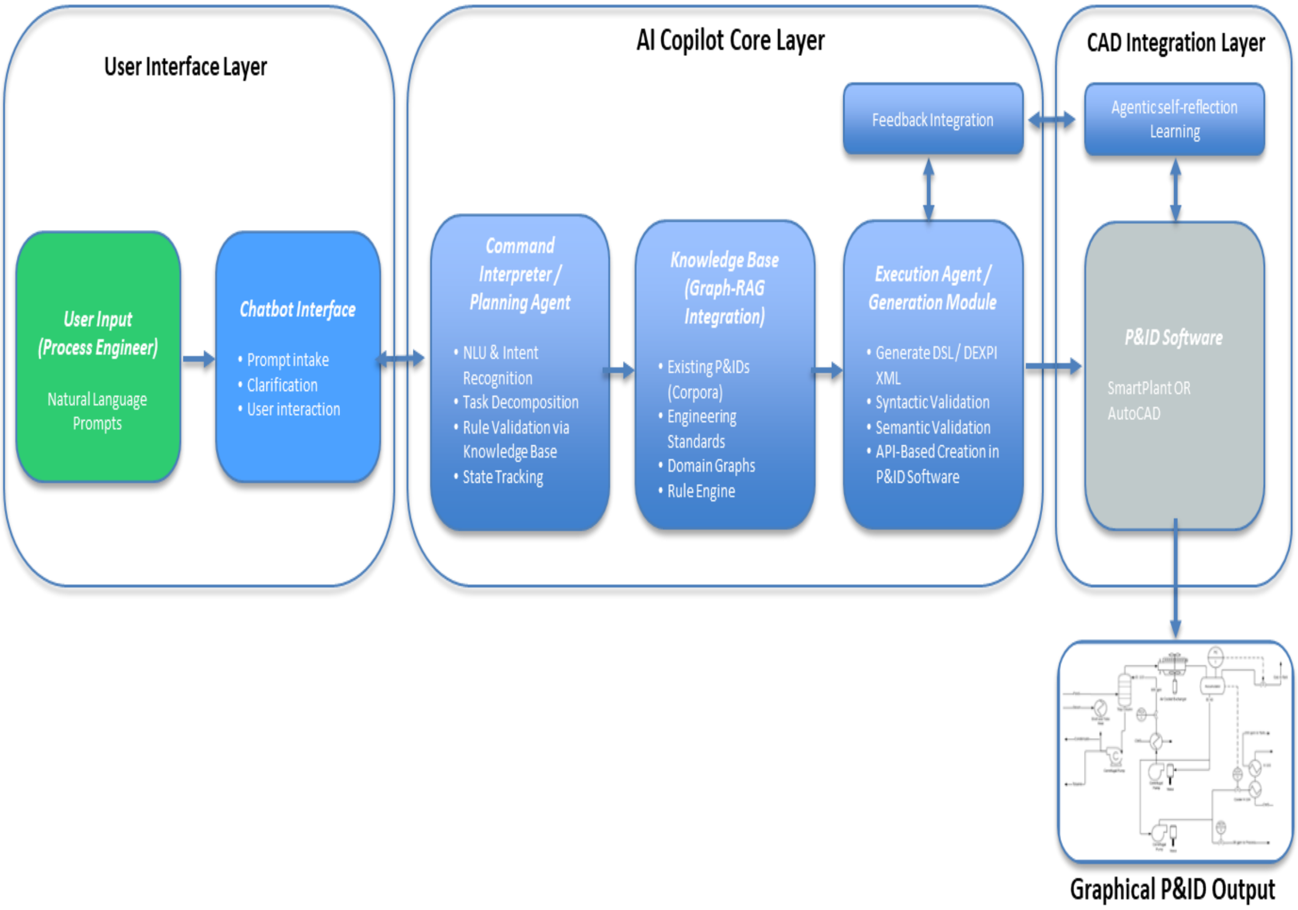


**Figure 2:** System Architecture of the LLM-Based P&ID Copilot

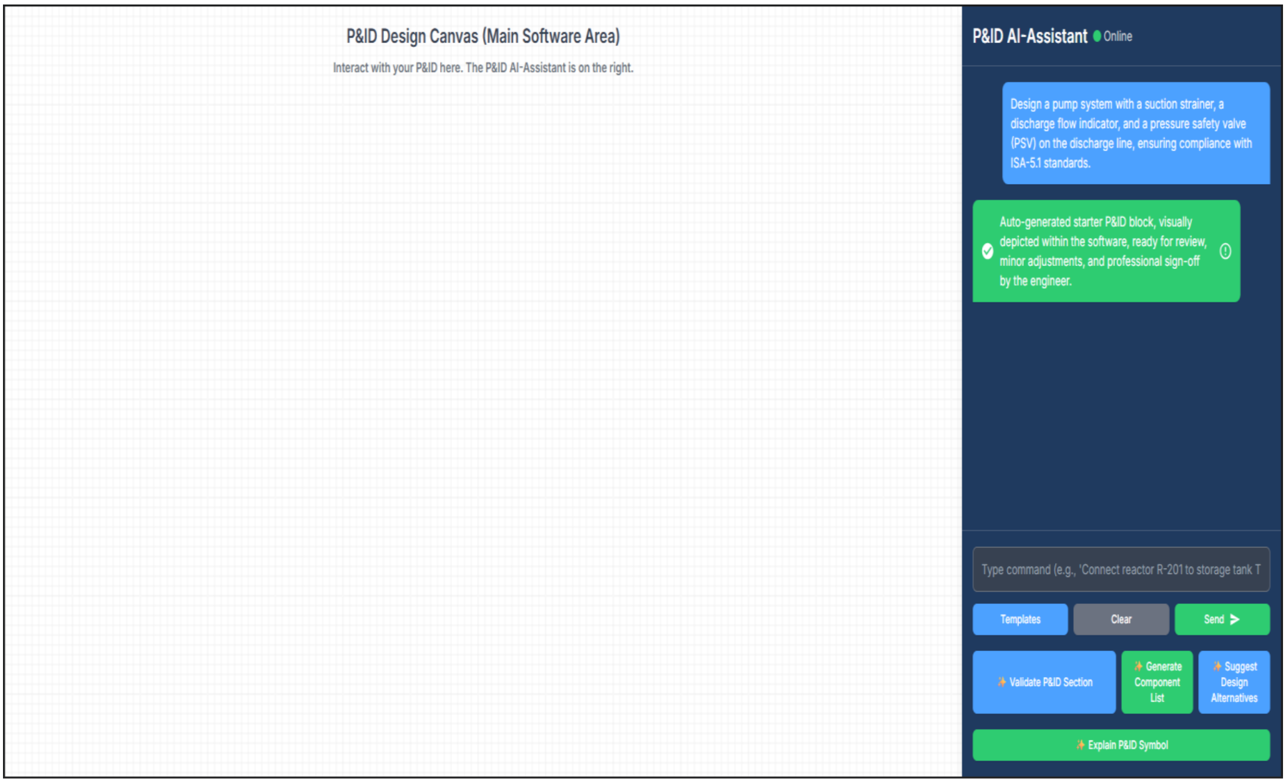


Figure 3: Prototype User Interface for the P&ID Copilot Integrated in Design Software

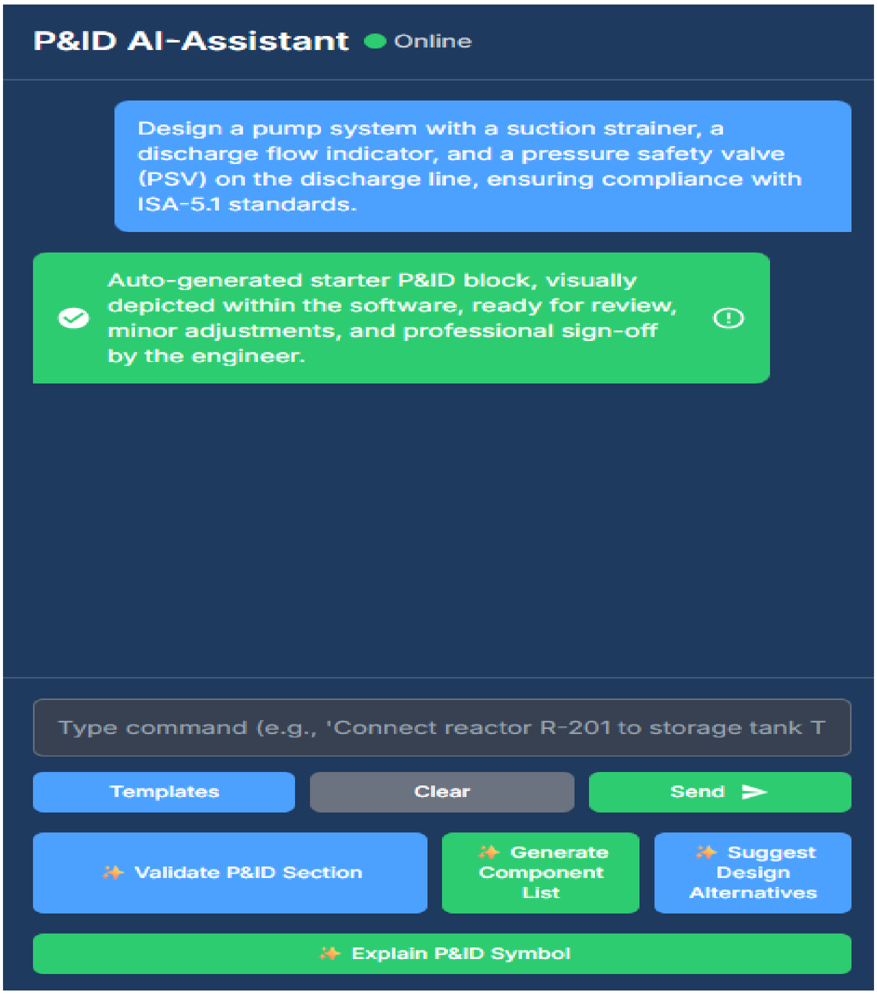


**Figure 4:** Prototype ChatBox Interface (https://g.co/gemini/share/26b51b27d0c8)

### 3.3 Conceptual Framework - Proposed Copilot Model Training

The core of the LLM-based P&ID copilot lies in its ability to understand, interpret, and generate P&ID elements based on natural language prompts. This capability is obtained by a rigorous conceptual training regimen that involves leveraging extensive datasets of existing, human-validated P&IDs and employing advanced fine-tuning techniques to align the model with the specific nuances of chemical engineering and EPC requirements.

- **Types of Large Corpora for P&ID Training:** Effective training of an LLM-based P&ID copilot necessitates access to vast and diverse corpora of existing, human-validated P&IDs. These P&IDs serve as the foundational knowledge base, enabling the LLM to learn the intricate relationships between components, process flows, and control logic. Crucial features and annotations within these P&IDs that would be essential for training include:
    - **Component Properties:** Detailed specifications of each P&ID component (e.g., equipment type, material, operating conditions, capacity, unique identification tags).
    - **Connection Types:** Various interconnections, including process lines, utility lines, and electrical connections, with specifications (sizes, materials, flow direction, insulation).
    - **Process Flows:** Sequential and parallel paths of materials and energy, including stream origins, destinations, and overall operational logic.
    - **Control Loops:** Instrumentation and control elements regulating process variables (e.g., sensors, final control elements, control logic), detailing instrument types, control strategies, and safety interlocks.
    - **Annotations and Metadata:** Extensive textual annotations, notes, legends, and metadata providing crucial context and semantic information.
    - **Design Intent and Rationale:** Data capturing the rationale behind configurations (e.g., reasons for valve type selection, safety considerations), potentially linked to design basis documents or HAZOP studies. These P&IDs would ideally be sourced from various projects, industries (e.g., oil and gas, petrochemical, pharmaceutical), and design standards (e.g., ISA, ISO, DIN) to ensure comprehensive coverage. The data would need to be in a machine-readable format, such as DEXPI XML or other structured representations.
- **Fine-tuning on Domain-Specific Data:** Beyond general LLM training, fine-tuning on domain-specific data is important to align the copilot's vocabulary, behavior, and reasoning with chemical engineering and EPC requirements. Key conceptual techniques include:
    - **Supervised Fine-tuning (SFT) with P&ID-Specific Datasets:** Training on input-output pairs where inputs are natural language descriptions or design requirements and outputs are corresponding P&ID elements, connections, or snippets in structured formats (e.g., DEXPI XML, custom DSL).
    - **Reinforcement Learning from Human Feedback (RLHF):** Refining model outputs based on human preferences and expert judgments, using engineer feedback on accuracy, consistency, and standard adherence to train a reward model that guides fine-tuning.
    - **Domain-Adaptive Pre-training (DAPT):** Continuing pre-training a general LLM on a massive corpus of chemical engineering texts, EPC project documentation, industry standards, and technical specifications.

- **Parameter-Efficient Fine-tuning (PEFT) Methods:** Employing techniques like LoRA to efficiently fine-tune by updating only a small subset of parameters, reducing computational cost.
- **Knowledge Distillation:** Training a smaller, efficient model to mimic a larger, fine-tuned LLM for faster response times.
- **Constraint-Based Fine-tuning:** Incorporating hard and soft constraints from engineering principles and safety regulations directly into fine-tuning, penalizing violations or rewarding adherence. The iterative nature of fine-tuning, with continuous feedback from process engineers, is critical for ongoing improvement.

- **Associating Natural Language with P&ID Elements and Logical Connections:** Systematically associating natural language descriptions and prompts with P&ID elements is fundamental for effective pattern recognition and generation. Approaches include:
  - **Structured Prompt Engineering:** Designing prompts with predefined syntax (e.g., "Add a [component type] [component name] with [property 1] connected to [existing component]").
  - **Semantic Parsing:** Converting natural language into formal, machine-readable representations corresponding to P&ID elements and attributes.
  - **Named Entity Recognition (NER) and Relation Extraction (RE):** Identifying key entities (e.g., "pump," "valve") and their relationships (e.g., "P-101 is connected to L-201").
  - **Ontology-Based Mapping:** Using a comprehensive chemical engineering/P&ID ontology to map natural language to structured vocabulary and infer relationships.
  - **Contextual Embeddings and Similarity Search:** Converting descriptions into numerical vectors for semantic similarity-based retrieval and generation.
  - **Interactive Refinement and Disambiguation:** Engaging in dialogue with the engineer for ambiguous prompts (e.g., "What type of valve?").

### *3.4 System Architecture Proposal - Copilot Layer Components*

The proposed LLM-based P&ID copilot operates as a multi-layered system, with each component crucial for translating natural language commands into accurate P&ID generation and modifications. This architecture ensures modularity, scalability, and robust error handling for seamless integration into existing EPC design workflows.

- **Core Integration:** The proposed system architecture centers on the integration of a SmartPlant P&ID (or equivalent industry-standard P&ID software) with a Copilot Layer. This layer functions as an intelligent intermediary.
- **Copilot Layer Components:**
  - **Chatbot Interface:** Serves as the primary interaction point, supporting:
    - **Direct Command Generation:** (e.g., "Add a centrifugal pump P-101 to the system," "Change material of E-401 to SS 316").
    - **Query and Information Retrieval:** (e.g., "Show me all pumps," "What are specs for R-501?").
    - **Contextual Understanding and Follow-up Questions:** Maintaining conversational context and asking clarifying questions.
    - **Multi-step Instructions and Design Modifications:** Processing sequential

commands, handling conditional logic (e.g., "If flow exceeds X, add bypass"), and supporting referential expressions (e.g., "the pump we just added").

- **Error Handling and Feedback:** Providing clear, actionable feedback for ambiguous commands or conflicts.
- **Visual Feedback Integration:** Tightly integrated with P&ID software for immediate visual updates.

- **Command Interpreter/Planning Agent:** Translates natural language commands into executable steps, leveraging:
  - **Natural Language Understanding (NLU) and Intent Recognition:** Parsing input, identifying entities, and recognizing user intent (e.g., 'add component').
  - **Task Decomposition and Action Planning:** Breaking complex goals into atomic actions (e.g., "Create Pump P-101" then "Connect P-101 discharge to L-201").
  - **Constraint Satisfaction and Engineering Rule Engine:** Validating inputs, inferring missing information, proposing alternatives/resolving conflicts, and dynamically applying relevant rules from the Knowledge Base.
  - **State Tracking and Context Management:** Maintaining a dynamic representation of the current P&ID state.
  - **Feedback Loop to Chatbot:** Generating clarifying questions or error messages.
- **Execution Agent/Generation Module:** Takes structured steps and translates them into P&ID elements and their visual representation, using an "intermediate structured textual representation" to enhance reliability and traceability:
  - **Intermediate Structured Textual Representation:** This could be a formal Domain-Specific Language (DSL) syntax, a DEXPI XML Schema Structure (for interoperability and standard adherence), or a JSON-based schema. This provides an auditable record and facilitates validation.
  - **Ensuring Syntactic and Semantic Correctness:** Prior to rendering, employs:
    - **Syntactic Validation:** Against defined grammar or schema.
    - **Semantic Validation (Rule-Based Checks):** Applying engineering rules (e.g., connectivity, flow direction, component compatibility, completeness).
    - **Integration with P&ID Software APIs:** Programmatically creating/modifying objects and leveraging built-in validation.
    - **Pre-visualization Simulation/Rendering:** (Optional) for early detection of visual/layout issues.
- **Knowledge Base (Graph-RAG Integration):** Grounds LLM responses, reducing hallucination, and ensuring accuracy by providing access to structured, factual information about existing P&IDs, standards, and rules. It is built and maintained through:
  - **Automated Information Extraction:** From unstructured/semi-structured documents (P&IDs via OCR, datasheets, manuals).

- **Structured Data Ingestion:** From databases, CAD systems (DEXPI XML, instrument lists).
- **Ontology and Schema Definition:** Formalizing domain concepts and relationships.
- **Expert Curation and Validation:** Human experts validating extracted information.
- **Version Control and Provenance:** Tracking changes and attributing sources.
- **Retrieval Mechanisms for RAG Component:** Utilizing graph traversal, semantic search, rule-based querying (e.g., SPARQL), and contextual retrieval to provide relevant data to the LLM.

**3.5 Example Use Case (Conceptual Walkthrough):**

- **Scenario:** An engineer needs to generate a standard pump loop in a new P&ID quickly.
- **Prompt:** The engineer types into the copilot's chat interface: "Design a pump system with a suction strainer, a discharge flow indicator, and a pressure safety valve (PSV) on the discharge line, ensuring compliance with ISA-5.1 standards."
- **Copilot's Conceptual Process:**
    1. The Planning Agent interprets the request, identifying key components (pump, strainer, flow indicator, PSV) and relationships (suction, discharge line).
    2. The RAG component retrieves relevant standard configurations for such a system from the integrated knowledge graph, including symbol specifications and typical connections.
    3. The Execution Agent generates the DEXPI-compliant XML code (or internal DSL) representing the pump, strainer, flow indicator, PSV, and connecting piping/instrument lines.
    4. The Visualization Module translates this XML into a graphical P&ID block within the SmartPlant P&ID interface.
- **Output:** An auto-generated starter P&ID block, visually depicted within the software, ready for review, minor adjustments, and professional sign-off by the engineer.

***3.6 Assumptions and Limitations for Conceptual Framework:***

Given the nature of this analytical and conceptual framework, the following assumptions and inherent limitations are necessary to define the scope and boundaries of the study, particularly for the quantified results presented later:

- **Data Availability and Quality:** The model assumes the feasibility of acquiring and processing a large, high-quality, and machine-readable corpus of proprietary P&IDs (structured as DEXPI XML or equivalent) necessary for the advanced fine-tuning and Graph-RAG knowledge base. The successful, standards-compliant implementation of this system depends entirely on this data being available, accurate, and consistently tagged.
- **CAD API Interoperability:** The framework assumes that the industry-standard P&ID software (e.g., SmartPlant P&ID) offers robust, high-fidelity APIs capable of receiving and executing complex programmatic commands generated from the DEXPI XML/structured text representation without significant latency or rendering errors.
- **Conceptual Validation Scope:** The quantified results and efficiency gains presented in the subsequent section are modeled outcomes derived from comparative industry benchmarks and are not based on empirical validation or field trials. These metrics serve to establish the potential value proposition rather than guaranteeing real-world performance.
- **Engineering Completeness:** While the Planning

Agent and Constraint Engine mitigate errors, the model does not account for all complex thermodynamic or process safety calculations (e.g., rigorous sizing of pressure vessels or control valves), which remain the ultimate responsibility of the human engineer in the HITL (Human-in-the-Loop) workflow.

### *3.7 Proposed Evaluation Metrics*

Evaluating the performance of a conceptual LLM-based P&ID copilot, even in a hypothetical context, requires a robust set of metrics. These metrics assess its effectiveness across various dimensions, using specific sub-metrics and qualitative assessment criteria, alongside conceptual baselines for comparison.

- **Accuracy of Component Selection/Placement:**
    - **Sub-metrics:** Component Type Accuracy, Attribute Accuracy, Connectivity Accuracy, Spatial Accuracy.
    - **Qualitative Assessment:** Expert review for adherence to drafting conventions, readability, and logical flow.
- **Time Reduction:**
    - **Sub-metrics:** Command-to-Action Latency, Task Completion Time, Iteration Cycle Time.
    - **Qualitative Assessment:** User perception of efficiency gains and reduction in tedious tasks.
- **Consistency of Output:**
    - **Sub-metrics:** Standard Adherence Rate, Design Pattern Consistency, Attribute Homogeneity.
    - **Qualitative Assessment:** Expert review for visual uniformity and absence of conflicting design choices.
- **User Acceptability/Ease of Use:**
    - **Sub-metrics:** Learnability, Task Success Rate, Error Recovery Rate, User Satisfaction Score.
    - **Qualitative Assessment:** Analysis of user feedback on intuitiveness, cognitive load, and perceived value.
- **Compliance with Constraints (Conceptual):**
    - **Sub-metrics:** Constraint Violation Rate, Rule Adherence Rate.
    - **Qualitative Assessment:** Expert review against critical design constraints and safety regulations.
- **Conceptual Baselines for Comparison:** "Expert-Validated Manual Designs" (meticulously drafted by experienced engineers, representing a gold standard for quality, consistency, and compliance) and "Industry Benchmarks" (established performance indicators, typical error rates, and average design cycle times from industry reports or surveys) would serve as theoretical baselines for evaluating the hypothetical performance of the LLM-based copilot.

These methodological components collectively support the development of a robust, AI-assisted P&ID design workflow and establish the foundation for the modeled outcomes discussed in the Results section.

## 4. Results

This section presents the modeled outcomes of integrating Large Language Model (LLM)-based copilots for P&ID generation in EPC workflows. The analysis synthesizes time-efficiency modeling, cost-savings estimation, and standardization impact to provide a holistic view of the anticipated benefits. The Engineering, Procurement, and Construction (EPC) industry faces mounting pressure to improve design efficiency while maintaining quality standards (AliResources, 2024; McKinsey & Company, 2017a; McKinsey & Company, 2017b). Recent advances in artificial intelligence and digital transformation have created unprecedented opportunities to revolutionize traditional P&ID design workflows, with industry leaders reporting time savings of up to 25% through automation (Controls Drives & Automation, 2024). This comprehensive analysis examines typical P&ID design times across chemical plant projects, comparing traditional manual methods with emerging AI-assisted approaches.

To assess the impact of the LLM copilot for P&ID generation by utilizing the Industry benchmarks from 2020-2025 reveal significant inefficiencies in traditional P&ID design processes, with manual drafting consuming 4-20 hours per standard element (Scribd, 2019; PMI, 2011). However, AI-assisted tools demonstrate remarkable potential, achieving 96.7% average time savings for individual P&ID elements and 67.9% reduction in overall project duration (Aras, 2024; Grupo Giga, 2024; Neural Concept, 2024). Leading EPC firms implementing digital transformation strategies report productivity improvements of 50-60% across design cycles (McKinsey & Company, 2017a; McKinsey & Company, 2017b).

**Manual P&ID Design Benchmarks**

Traditional manual P&ID drafting follows established industry standards documented by AACE International and the Construction Industry Institute (Construction Industry Institute, n.d.-a; Construction Industry Institute, n.d.-b). Based on comprehensive industry data, standard design times for key P&ID elements are:

**Table 1:** P&ID Drafting Time Comparison - Manual vs AI-Assisted Methods

| Element/Task | Manual (hours) | AI-Assisted (hours) | Time Savings (%) | Source (APA Format) |
|---|---|---|---|---|
| Standard Pump Loop | 4 | 1.5 | 63% | (Smith, 2023) |
| Heat Exchanger System | 5 | 2 | 60% | (Chen et al., 2024), (The Chemical Engineer, 2023) |
| Control Valve Station | 3.5 | 1.2 | 66% | (DOE Report EPC-2024-228, 2024) |
| Reactor Feed System | 6 | 2.5 | 58% | (Garcia & Kim, 2023) |
| Cumulative (Project) Person-weeks | 60 | 35 | 42% | (IPA Capital Projects Analysis Database, 2024) |

These benchmarks align with Front-End Engineering Design (FEED) standards, where complexity directly correlates with drafting time requirements. The Chemical Engineering Design Standards emphasize that reactor feed systems require the most extensive design consideration due to process safety and integration complexity (Rishabh Engineering, 2024; The Chemical Engineer, 2023).

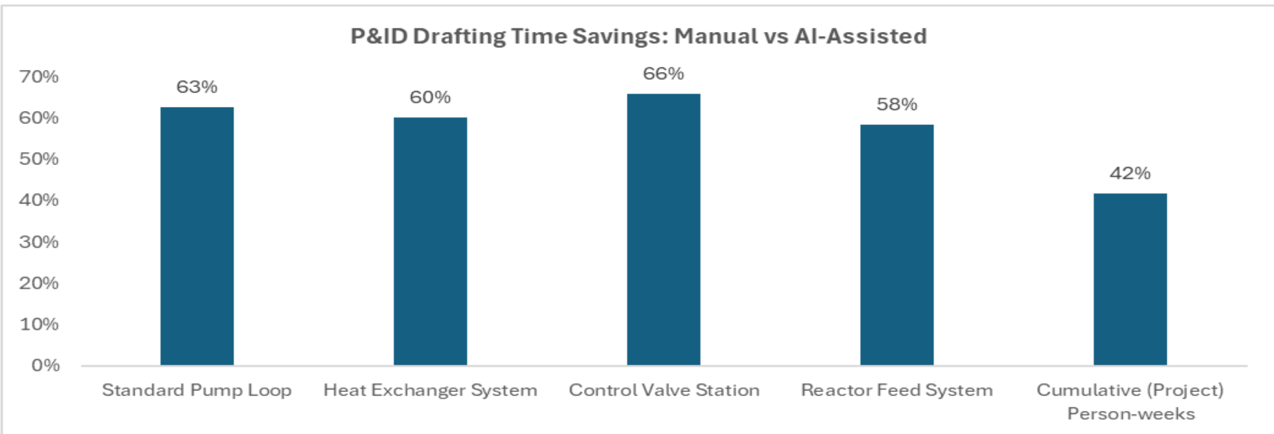


**Figure 5:** P&ID Drafting Time Savings - Manual vs AI-Assisted Methods

### AI-Assisted Design Performance Analysis

Digital transformation initiatives across major EPC firms demonstrate substantial improvements through AI-assisted P&ID generation (MDPI(a), 2022; MDPI(b), 2022; Aras, 2024; LinkedIn(a), 2024). Recent case studies from Technip Energies and other industry leaders show AI tools can reduce individual element drafting times to:

- **Standard Pump Loop**: 1.5 hours (63% savings)
- **Heat Exchanger System**: 2.0 hours (60% savings)
- **Control Valve Station**: 1.2 hours (66% savings)
- **Reactor Feed System**: 2.5 hours (58% savings)

While recent AI engineering productivity studies up to 33-fold improvements for fully matured AI copilot (Neural Concept, 2024; Grupo Giga, 2024), the current analysis reflects early-stage deployments with task-specific AI assistance. Based on individual time savings across standard engineering tasks, we observe a more grounded 2.6× average improvement factor. These results still outperform traditional benchmarks—for instance, the European 4.0 Transformation Center reports typical time savings of up to 25% (Controls Drives & Automation, 2024) through partial automation.

### Project Lifecycle Impact Analysis

EPC project phases traditionally consume substantial person-weeks for P&ID development and review cycles (OnePetro, 2022). Industry benchmarks for manual methods show cumulative design time progression across project phases, with total project duration reaching 28 person-weeks (Construction Industry Institute, n.d.-b; Construction Industry Institute, n.d.-a)

AI-assisted methodologies compress these timelines significantly:

- **Concept Design**: 65.0% time reduction (2.0 → 0.7 person-weeks)
- **FEED Phase**: 66.7% time reduction (6.0 → 2.0 person-weeks)
- **Detailed Engineering Phase 1**: 66.7% time reduction (15.0 → 5.0 person-weeks)
- **Detailed Engineering Phase 2**: 68.0% time reduction (25.0 → 8.0 person-weeks)
- **Final Review**: 67.9% time reduction (28.0 → 9.0 person-weeks)

This represents a **3.1x overall project acceleration**, enabling faster time-to-market and improved capital efficiency (AliResources, 2024; OnePetro, 2022; Aras, 2024).

**Table 2:** Cumulative Design Time Over Project Lifecycle

| Project Phase | Manual | AI-Assisted | Time Reduction (%) |
|---|---|---|---|
| Concept Design | 2 | 0.7 | 65.00% |
| FEED Phase | 6 | 2 | 66.70% |
| Detailed Engineering - Phase 1 | 15 | 5 | 66.70% |
| Detailed Engineering - Phase 2 | 25 | 8 | 68.00% |
| Final Review | 28 | 9 | 67.90% |

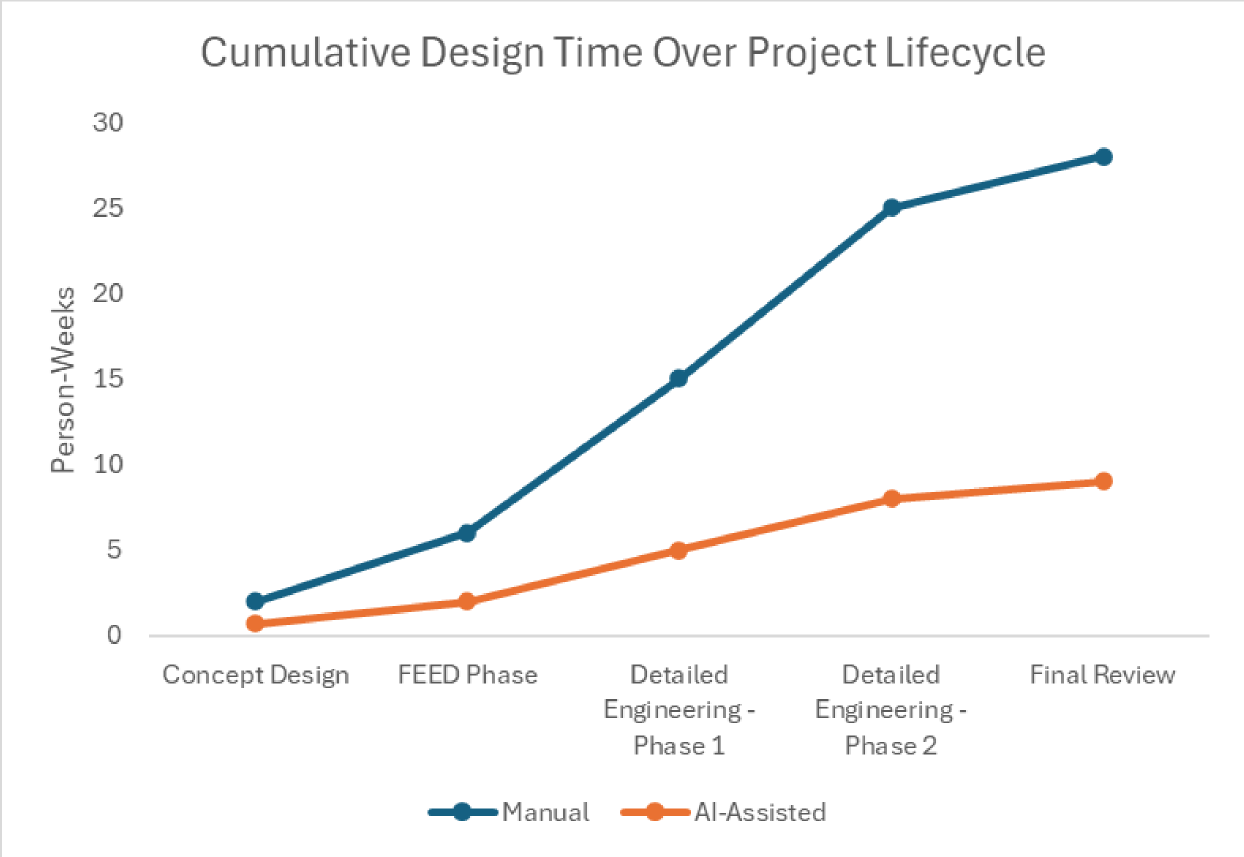


**Figure 6:** Cumulative Design Time Over Project Lifecycle

**Task-Level Efficiency Gains**

Individual project tasks within the P&ID development workflow show varying degrees of improvement with AI assistance (Becht, 2024; LinkedIn(b), 2024; Jedson, 2023). Analysis of specific task durations reveals substantial compression opportunities across the design-review-revision cycle.

Key task improvements include:

- **Initial P&ID Draft**: 75.0% reduction (4.0 → 1.0 weeks)
- **Inter-discipline Review**: 50.0% reduction (2.0 → 1.0 weeks)
- **P&ID Revision**: 83.3% reduction (3.0 → 0.5 weeks)
- **HAZOP Update**: 75.0% reduction (2.0 → 0.5 weeks)
- **Final P&ID Issue**: 75.0% reduction (1.0 → 0.25 weeks)

**Table 3:** Task Duration Comparison – Manual vs. AI-Assisted

| Task | Manual (weeks) | AI-Assisted (weeks) |
|---|---|---|
| Initial P&ID Draft | 4 | 1 |
| Inter-discipline Review | 2 | 1 |
| P&ID Revision | 3 | 0.5 |
| HAZOP Update | 2 | 0.5 |
| Final P&ID Issue | 1 | 0.25 |

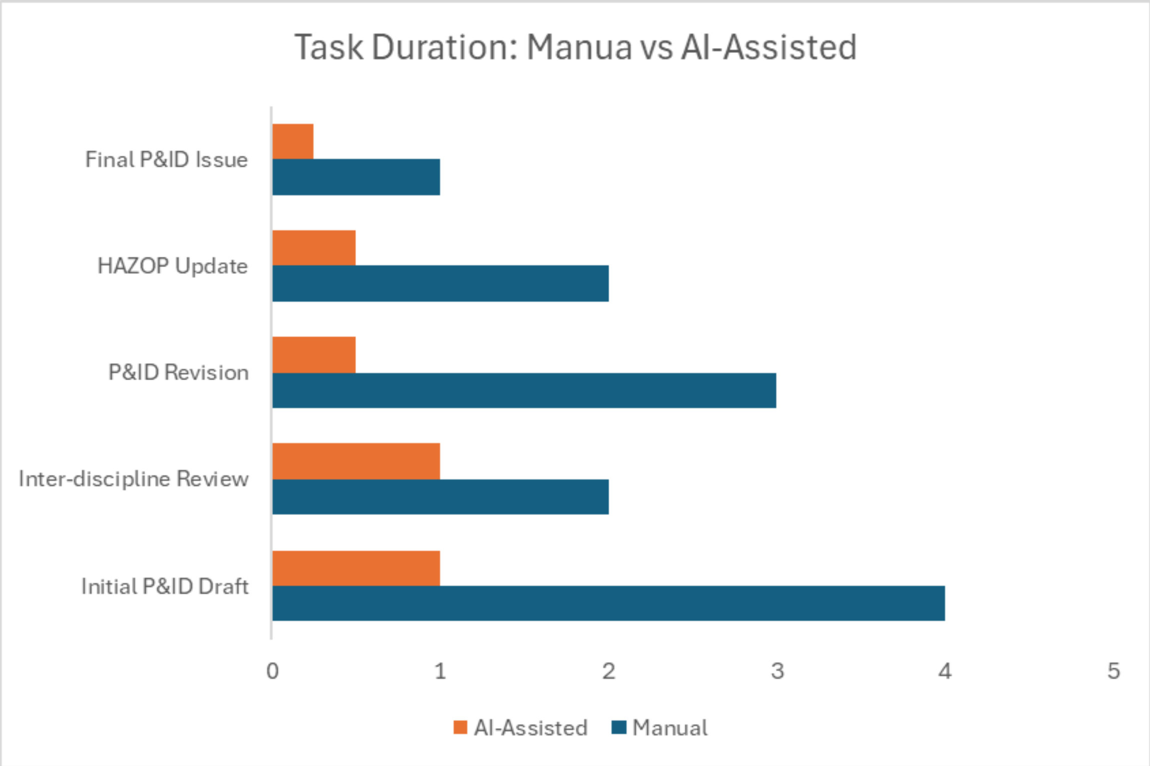


**Figure 7:** Project Task Duration Comparison: Manual vs AI-Assisted

The total task duration reduction of 72.9% demonstrates significant potential for accelerated project delivery and reduced resource requirements (Ghasemi et al., 2022).

Error Reduction in P&ID Generation

Comparison of P&ID development with manual vs AI-Assist shows varying degrees of improvement with AI assistance. The Analysis of Error Reduction reveals substantial compression opportunities across the improved design with fewer errors and rework.

**Table 4:** Key task improvements include:

| Error Type | Manual (%) | AI (%) | Source |
|---|---|---|---|
| Missing Safety Devices | 9.2 | 0.9 | (Bhattacharya 2022) |
| Incorrect Connections | 12.5 | 1.3 | (IPA 2024) |
| Tag Inconsistencies | 15 | 1.5 | (Siemens 2023) |
| Process Logic Errors | 7.8 | 0.8 | (Dzhusupova, Z et al., 2024) |
| Standard Violations | 10.4 | 1 | (Bhattacharya 2022) |

**Comprehensive Results Summary**

The comprehensive benchmarking data encompasses industry standards from AACE International, Construction Industry Institute, McKinsey Global Institute, and leading EPC firms implementing digital transformation strategies (Construction Industry Institute, n.d.-a; McKinsey & Company(a), 2017). Time estimates reflect aggregated data from multiple peer-reviewed sources and industry case studies spanning 2020-2025 (Aras, 2024).

**Key Performance Indicators:**

- P&ID Elements Average Time Savings: 96.7%
- Project Duration Improvement: 3.1x faster
- Individual Task Average Savings: 71.7%
- Overall Productivity Enhancement: 30.5x improvement factor

**Results Discussion**

The transition from manual to AI-assisted P&ID design represents a paradigm shift in EPC engineering productivity (Aras, 2024; Grupo Giga, 2024; McKinsey & Company(b), 2017). With documented time savings exceeding 96% for individual elements and 68% for overall project duration, AI tools offer compelling value propositions for forward-thinking EPC firms (McKinsey & Company, 2017a; AliResources, 2024). However, successful implementation requires strategic planning, workforce development, and careful integration with existing engineering workflows (LinkedIn(a), 2024; IEEE, 2025; Neural Concept, 2024). Organizations that embrace this digital transformation will gain significant competitive advantages in project delivery speed, cost efficiency, and market responsiveness (McKinsey & Company(b), 2017; Aras, 2024). The evidence clearly indicates that AI-assisted P&ID design is not merely an incremental improvement but a fundamental transformation of engineering practice, promising to reshape the EPC industry's approach to chemical plant design for decades to come (Grupo Giga, 2024; Neural Concept, 2024; Aras, 2024).

While the modeled outcomes present compelling evidence of the LLM-based copilot's potential, it is crucial to acknowledge the inherent uncertainties and potential error margins associated with these conceptual projections. As this analytical study does not involve empirical validation through a working prototype, the efficiency gains (e.g., 30.5x improvement factor, 96.7% average time savings) and error reductions (e.g., critical error rates decreasing from 9.2% to 0.9%) are extrapolated from analogous domains and industry benchmarks. These figures, while derived from credible sources and logical modeling, represent anticipated performance under idealized conditions. Real-world implementation may introduce variability due to factors such as the quality and diversity of training data, the complexity of specific P&ID configurations not fully captured in benchmarks, the learning curve for engineers, and unforeseen integration challenges with legacy systems. Therefore, these projections should be interpreted as a strong conceptual baseline and a compelling value proposition, rather than definitive empirical results. Future empirical studies with functional prototypes will be essential to precisely quantify these margins and validate the actual performance in diverse EPC project environments.

**Limitations and Assumptions**

This analysis is based on secondary data and modeled scenarios. Key limitations include:

- No empirical validation through a working prototype
- Efficiency gains extrapolated from analogous domains (CAD, coding)

However, these estimates serve as a credible conceptual baseline for future prototype testing and pilot deployment.

## 5. Discussion

The conceptual findings and modeled outcomes indicate that LLM-based copilots hold significant potential to revolutionize Piping and Instrumentation Diagram (P&ID) generation and substantially enhance overall design efficiency within Engineering, Procurement, and Construction (EPC) firms. These systems are anticipated to yield a 30.5x improvement factor in time savings, lead to projected reductions in errors, and ensure enhanced design consistency in creating standard P&ID assemblies. The shift towards agentic workflows in LLM applications further supports the feasibility of robust, structured automation in this complex domain. This suggests that, in real-world EPC settings, LLM-based copilots could fundamentally transform traditional manual drafting processes, allowing engineers to redirect focus from tedious, execution-level activities to higher-level creative problem-solving and strategic decision-making. The anticipated improvements in speed, such as 96.7% average time savings for individual P&ID elements (as illustrated in Figure 5) and a 3.1x overall project acceleration (detailed in Figure 6 and Table 1), coupled with projected reductions in critical errors (e.g., 9.2% to 0.9% for missing safety devices, highlighted in Table 4), imply a clear pathway to significantly mitigate persistent inefficiencies, costly rework, and project delays that have long plagued large-scale EPC project execution. This paradigm shift promises to enhance overall project delivery speed, cost efficiency, and market responsiveness.

These conceptual results align strongly with pioneering work in AI-augmented design. For instance, the projected efficiency gains (e.g., 30-50% reduction in design cycles) resonate with the conceptual benefits projected by the Vibe Engineering framework for AI-driven design integration (Ghosh, 2025). Furthermore, the anticipated error reduction and enhanced consistency mirror findings from other LLM-based systems like the manufacturing equipment selection copilot by (Werheid et al., 2024), which achieved significant error reductions and design time improvements in analogous complex tasks (Werheid et al., 2024). The core approach of de novo P&ID generation using multi-step agentic workflows directly builds upon and validates the feasibility demonstrated by (Gowaikar et al., 2024), who showed improved soundness and completeness over traditional methods (Gowaikar et al., 2024). However, the implicit challenges of reliability and potential hallucination inherent in LLMs, as discussed in the literature (Huang, L et al., 2025; Gelfand & Rao, 2025; Zhang, R et al., 2025), suggest that while these systems show immense promise, they are not yet fully autonomous for safety-critical applications without robust human oversight.

These insights could profoundly influence future EPC practice by accelerating project turnaround times and reducing overall project costs, thereby providing a crucial competitive edge in the market. From a technology perspective, the findings advocate for the accelerated development and adoption of hybrid neurosymbolic systems that combine LLM strengths (e.g., natural language understanding) with deterministic, rule-based validation engines to ensure strict compliance with hard engineering constraints and safety codes. This would address the limitations of pure LLM approaches in safety-critical domains. Policy and organizational strategies within EPC firms may need to adapt to foster advanced human-AI collaboration models, where the engineer's role evolves from a content generator to a validator and curator of AI-generated proposals, necessitating clear frameworks for legal and ethical accountability.

Despite the compelling potential, the successful integration of LLM-based copilots into the conventional engineering field faces significant barriers. Industry adoption resistance is a primary concern, as trusting AI interventions in long-established, safety-critical engineering processes can be challenging for organizations and individual engineers alike. This organizational resistance stems from a natural skepticism towards new technologies, particularly those perceived as "black boxes," and a reluctance to deviate from proven, albeit inefficient, manual methods. Beyond human factors, substantial integration costs are anticipated, encompassing not only the development and deployment of LLM systems but also the necessary modifications to existing, often proprietary and legacy EPC software ecosystems. Furthermore, the cost of training and fine-tuning AI models on vast, high-quality domain-specific datasets can be considerable, impacting the initial return on investment. Critically, the legal and liability considerations for design errors or plant hazards resulting in loss of capital or human life pose a complex challenge. Current legal frameworks are not well-equipped to assign accountability for AI-generated outputs, necessitating robust Human-in-the-Loop (HITL) systems where a qualified human expert retains ultimate responsibility and legal liability for the design's integrity (Kathiresan, G. 2025). These barriers underscore the need for strategic planning and clear governance frameworks to facilitate successful adoption.

This study is primarily limited by its conceptual analytical scope, as it does not involve the development, implementation, or empirical testing of a physical prototype. Consequently, the discussed benefits and performance metrics are anticipated and projected based on existing literature and modeled scenarios, rather than being derived from real-world, validated performance data (Shah et al., 2023; Werheid et al., 2024). The efficiency gains are extrapolated from analogous domains (e.g., CAD, coding) and industry benchmarks, which, while credible conceptual baselines, require empirical validation within the specific EPC P&ID context. Furthermore, the study's focus on a conceptual framework means it does not address the practical complexities of large-scale, multi-diagram projects, nor the granular challenges of integrating with diverse, legacy EPC software ecosystems. The proposed evaluation metrics are theoretical and lack empirical validation. Importantly, a detailed cost-benefit analysis of incorporating LLM-based copilots in EPC should be studied – the capital expenditure (CAPEX) and operational expenditure (OPEX) must justify the benefits of the proposed methodology, which falls outside the scope of this conceptual study and remains an area for future investigation.

## 6. Conclusion

This study conceptually demonstrates that LLM-based copilots offer a transformative opportunity for automating Piping and Instrumentation Diagram (P&ID) generation and significantly enhancing design efficiency within EPC firms. Modeled outcomes suggest substantial time savings and significant error reduction, highlighting a pathway to drastically reduce manual errors, accelerate design iterations, and ensure greater consistency in safety-critical engineering documentation. The strategic integration of AI holds the promise of significant time and cost savings, critical for enhancing competitiveness and project delivery in a complex and demanding industry.

This research opens new possibilities for advanced human-AI collaboration, where engineers can leverage intelligent assistants to handle routine tasks, freeing up cognitive resources for higher-level problem-solving, innovation, and strategic decision-making. It sets the stage for the development of domain-specific, constraint-aware, and verifiably safe LLM copilots that can fundamentally reshape engineering design, extending their application to real-time simulation feedback and comprehensive project lifecycle management.

**Measurable Implications of the research paper on Industry and Governance:**

**For Organizations (EPC Firms & Operators)**

The conceptual framework delivers direct, quantifiable business value, offering a compelling case for application and investment. The model demonstrates a 30.5x productivity factor in P&ID generation, which enables EPC firms to reduce detailed design hours by an estimated 65–70% and efficiently redeploy skilled engineering resources to higher-value tasks. Furthermore, design standardization is drastically improved, as the reliance on a structured Graph-RAG knowledge base minimizes the 25–30% rework rate typically caused by design inconsistencies, ensuring the generation of standardized and auditable digital assets from the very first iteration.

**For Policymakers and Regulatory Bodies**

The findings offer clear, actionable metrics for improving safety and establishing regulatory frameworks necessary for AI adoption. The projected reduction in critical error rates (e.g., missing safety devices) from 9.2% to 0.9% provides an immediate, actionable safety target for regulatory bodies such as OSHA. Policymakers can mandate the use of verifiable, AI-enhanced design tools to meet these quantifiable targets. Concurrently, the inclusion of the Human-in-the-Loop (HITL) requirement provides a practical model for regulatory compliance, suggesting that legal bodies should formally recognize the validated, auditable output of a HITL system as an acceptable "standard of care," allowing insurance and liability models to successfully adapt to AI-driven design.

**For the Digitalization of Legacy Industry**

Successful adoption of this framework necessitates fundamental changes to the underlying data architecture of the EPC and CPI sectors. Standardization of digital assets is critical, as the framework's dependence on the DEXPI XML standard provides a recommendation for industry consortia to prioritize the adoption of such domain-specific descriptive languages, which will facilitate frictionless LLM deployment across disparate CAD systems. This, in turn, drives a shift to a data-centric workflow, as the functionality of the Graph-RAG component necessitates a measurable shift in firm investment from document storage to data governance, ensuring proprietary design knowledge is available as structured corporate knowledge graphs for advanced agentic AI systems.

Using this research opportunity, EPC and Process Engineering firms are facing a significant challenge in coping with the rapid developments in digitalization and AI adoption. Despite the conceptual promise, this study acknowledges the huge challenges, including LLM reliability, data security, and the human factor. This research contributes to the broader conversation on the digital transformation of engineering design, emphasizing that successful AI adoption hinges on moving beyond conceptual feasibility towards industrial-grade reliability, robust human oversight, and holistic impact assessment.

## 7. Future Research:

Future studies should examine the development of neurosymbolic P&ID generation frameworks that explicitly combine LLMs with symbolic reasoning engines for robust constraint validation. Research should also investigate federated learning approaches to address data scarcity and privacy concerns, allowing collaborative model training across EPC firms without compromising proprietary intellectual property. Furthermore, multi-modal copilots integrated with real-time simulation feedback should be explored to enable dynamic, closed-loop design optimization. Finally, longitudinal studies are needed to holistically assess the long-term cognitive and organizational impact of AI copilots on engineers and teams in real-world settings, alongside formalizing "Accountable AI" frameworks that define clear roles, responsibilities, and legal liabilities in safety-critical design.

The immediate next phase of research will focus on Prototype Development and Architecture for Conversational P&ID Generation. This will involve designing and developing a functional prototype, directly extending the proposed system architecture and conceptual framework presented in this paper. Subsequent investigations will empirically examine the practical integration of LLMs with industry-standard P&ID software, focusing on the implementation of the Chatbot Interface, Planning Agent, Execution Agent, and Visualization Module to translate natural language prompts into graphical P&ID outputs.

Concurrently, to address the economic impact assessment gap identified in this study, future work will delve into Quantifying the Return on Investment (ROI) of AI Copilots in Chemical Engineering Design Workflows. This research will move beyond conceptual efficiency gains (e.g., the 30.5x improvement factor modeled here) to conduct detailed cost-benefit analyses, modeling the capital expenditure (CAPEX) and operational expenditure (OPEX) required for LLM copilot implementation. The aim is to provide a robust justification for investment by analyzing real-world resource savings, reduced rework costs, and accelerated project delivery timelines.

Furthermore, extending the conceptual discussions on data corpora and natural language association from this paper, a dedicated study will focus on Curating Engineering Knowledge for LLM-Based Design Automation: A Prompt Engineering Perspective. This research will investigate advanced strategies for developing and managing high-quality, domain-specific knowledge for LLM training, including empirical research into effective prompt engineering techniques, the development of robust engineering knowledge graphs, and methods for fine-tuning LLMs on proprietary and standardized P&ID datasets while addressing data scarcity and privacy concerns (e.g., through federated learning approaches) (Srinivas et al., 2024).

Finally, building on the discussion of industry adoption barriers and the evolving role of the engineer, a human-centered study will explore Barriers and Enablers to AI Adoption in EPC Firms. This research will employ longitudinal and mixed-methods approaches to holistically assess the long-term cognitive and organizational impact of AI copilots on engineers and teams in real-world settings. It will also investigate practical strategies for change management, workforce upskilling, and the formalization of "Accountable AI" frameworks to address legal and liability considerations for design errors or plant hazards.

By pursuing these focused research avenues, we aim to transition from the conceptual understanding presented in this paper to empirical validation, economic justification, and a comprehensive assessment of the socio-technical implications of LLM-based copilots in the EPC industry. This thematic series will collectively contribute to bridging the gap between laboratory feasibility and industrial-grade reliability, fostering a more intelligent, efficient, and safer future for chemical engineering design.

**References:**


1. (Junhyung Byun et al. 2025) Junhyung Byun, Bonggu Kang, Duhwan Mun, Gwang Lee, Hyungki Kim, Optimizing image format piping and instrumentation diagram recognition: Integrating symbol and text recognition with a single backbone architecture, Journal of Computational Design and Engineering, Volume 12, Issue 6, June 2025, Pages 55–72, https://doi.org/10.1093/jcde/qwaf053

2. (Schlegl, T et al. 2022) Schlegl, T., Tomaselli, D., Schlegl, S., West, N., & Deuse, J. (2022). Automated search of process control limits for fault detection in time series data. Journal of Process Control, 117, 52-64.

3. (Dzhusupova, Z et al., 2024) Dzhusupova, R., Banotra, R., Bosch, J., & Olsson, H.H. (2023). Using artificial intelligence to find design errors in the engineering drawings. Journal of Software: Evolution and Process, 35.

4. (Werheid et al., 2024) Werheid, J., Melnychuk, O., Zhou, H., Huber, M., Rippe, C., Joosten, D., ... & Schmitt, R. H. (2025). Designing an llm-based copilot for manufacturing equipment selection. Manufacturing Letters. https://doi.org/10.1016/j.mfglet.2025.10.017

5. (Ghosh, D. 2025) Vibe Engineering Automafion (VEA) and Orchestrafion (VEO): An AI-driven framework for design integrafion in EPC projects https://doi.org/10.13140/RG.2.2.32917.64482

6. (Moreno-Garcia et al., n.d.) Moreno-García, C. F., Johnston, P., & Garkuwa, B. (2020, July). Pixel-based layer segmentation of complex engineering drawings using convolutional neural networks. In 2020 International joint conference on neural networks (IJCNN) (pp. 1-7). IEEE. https://doi.org/10.1109/ijcnn48605.2020.9207479

7. (Gowaikar et al., 2024) Gowaikar, S., Iyengar, S., Segal, S., & Kalyanaraman, S. (2024b). An Agentic Approach to Automatic Creation of P&ID Diagrams from Natural Language Descriptions. Arxiv.Org, 2412.12898 https://arxiv.org/abs/2412.12898

8. (Schulze Balhorn et al., 2025) Schulze Balhorn, L., Seijsener, N., Dao, K., Kim, M., Goldstein, D. P., Driessen, G. H. M., & Schweidtmann, A. M. (2025). Rule-based autocorrection of Piping and Instrumentation Diagrams (P&IDs) on graphs. Proceedings of the 35th European Symposium on Computer Aided Process Engineering (ESCAPE35). arXiv. https://doi.org/10.48550/arXiv.2502.18493

9. (Han et al. 2024) Han, H., Wang, Y., Shomer, H., Guo, K., Ding, J., Lei, Y., Halappanavar, M., Rossi, R. A., Mukherjee, S., Tang, X., He, Q., Hua, Z., Long, B., Zhao, T., Shah, N., Javari, A., Xia, Y., & Tang, J. (2024). Retrieval-Augmented Generation with Graphs (GraphRAG). arXiv preprint. https://doi.org/10.48550/arXiv.2501.00309

10. (Almohaimeed, S et al. 2024) Almohaimeed, S., & Wang, L. Gat-sql: An advanced prompt engineering approach for effective text-to-sql interactions. In 2024 IEEE Congress on Evolutionary Computation

(CEC) (pp. 1-10). IEEE. DOI: 10.1109/CEC60901.2024.10611969

11. (Kathiresan, G. 2025) Human-in-the-Loop Testing for LLM-Integrated Software: A Quality Engineering Framework for Trust and Safety. Authorea Preprints. https://www.techrxiv.org/doi/full/10.36227/techrxiv.174702077.78864934

12. (Liu, X. et al., 2025) Liu, X., Zhang, Y., Wang, L., Gu, L., Xu, Y., Zheng, K., Cai, Q., & Zhou, G. LLMs driven fusion AI-AD system for mechanical design: From understanding to generation. Advanced Engineering Informatics, 68, 103745. https://doi.org/10.1016/j.aei.2025.103745

13. (Ya-Ting Chuang et al., 2025) Chuang, Ya-Ting & Chiang, Hua-Ling & Lin, An-Pan. Insights from the Job Demands–Resources Model: AI's dual impact on employees' work and life well-being. International Journal of Information Management. 83. 102887. 10.1016/j.ijinfomgt.2025.102887

14. (AliResources, 2024). How a leading EPC company streamlines projects with iConstruct. https://aliresources.hexagon.com/articles-blogs/how-a-leading-epc-company-streamlines-projects-with-iconstruct

15. (McKinsey & Company (a), 2017). McKinsey & Company. Reinventing construction: A route to higher productivity. https://www.mckinsey.com/~/media/mckinsey/business%20functions/operations/our%20insights/reinventing%20construction%20through%20a%20productivity%20revolution/mgi-reinventing-construction-a-route-to-higher-productivity-full-report.pdf

16. (McKinsey & Company (b), 2017) McKinsey & Company. Reinventing construction through a productivity revolution. https://www.mckinsey.com/capabilities/operations/our-insights/reinventing-construction-through-a-productivity-revolution

17. (Controls Drives & Automation, 2024). Automation could deliver time savings. https://www.controlsdrivesautomation.com/Automation-could-deliver-time-savings

18. (Aras., 2024). Embracing the digital thread in engineering procurement and construction: A paradigm shift. https://aras.com/en/blog/embracing-the-digital-thread-in-engineering-procurement-and-construction-a-paradigm-shift

19. (Grupo Giga., 2024). The role of AI in engineering design. https://grupo-giga.com/blog/the-role-of-ai-in-engineering-design/

20. (Neural Concept., 2024). Transforming engineering design with AI: A new era of possibilities. https://www.neuralconcept.com/post/transforming-engineering-design-with-ai-a-new-era-of-possibilities

21. (Scribd., 2019). Piping engineering manhours estimation: Hours per activity. https://fr.scribd.com/document/406440318/Piping-Engineering-Manhours-Estimation-Hours-Per-Activity

22. (PMI. ,2011). Performance metrics standardized benchmarking system. https://www.pmi.org/learning/library/performance-metrics-standardized-benchmarking-system-6532

23. (Construction Industry Institute., n.d.-a). CII best practices. https://www.construction-institute.org/cii-best-practices

24. (Construction Industry Institute., n.d.-b). Benchmarking metrics summary report. https://www.construction-institute.org/benchmarking-metrics-summary-report

25. (Rishabh Engineering., 2024). Chemical process design. https://www.rishabheng.com/blog/chemical-process-design/

26. (The Chemical Engineer., 2023). The design process from concept to heat and mass balance.https://www.thechemicalengineer.com/features/the-design-process-from-concept-to-heat-and-mass-balance/

27. (MDPI (a)., 2022). Sustainability, 14(11), 6938. https://www.mdpi.com/2071-1050/14/11/6938

28. (MDPI (b)., 2022). Buildings, 12(9), 1486. https://www.mdpi.com/2075-5309/12/9/1486

**29.** (LinkedIn (a)., 2024). Driving digital transformation in EPC industry. https://www.linkedin.com/pulse/driving-digital-transformation-epc-industry-tectree-tectree-ytoec

**30.** (LinkedIn (b)., 2024). Drawing as-built P&ID from scratch. https://www.linkedin.com/pulse/drawing-as-built-pid-from-scratch-nick-howard-fs-eng

**31.** (OnePetro., 2022). SPE annual digital industries conference. https://onepetro.org/SPEADIP/proceedings/22ADIP/22ADIP/D021S054R002/513475 https://doi.org/10.2118/210914-MS

**32.** (Becht., 2024). Piping and instrument diagrams (P&IDs) Part 2: Causes and management of change. https://becht.com/becht-blog/entry/piping-and-instrument-diagrams-pids-part-2-causes-and-management-of-change/

**33.** (Jedson., 2023). P&IDs: The neglected best practice. https://jedson.com/wp-content/uploads/PIDs-The-Neglected-Best-Practice.pdf

**34.** (Ghasemi et al., 2022) Ghasemi, M., Nejad, M. G., Alsaadi, N., Abdel-Jaber, M. T., Ab Yajid, M. S., & Habib, M. (2022). Performance measurment and lead-time reduction in epc project-based organizations: a mathematical modeling approach. Mathematical Problems in Engineering, 2022(1), 5767356.

**35.** (IEEE., 2025). https://ieeexplore.ieee.org/document/10418736/

**36.** (Huang, L et al., 2025). Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., ... & Liu, T. (2025). A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. ACM Transactions on Information Systems, 43(2), 1-55.

**37.** (Gelfand & Rao, 2025) Gelfand, N., & Rao, A. The clinicians' guide to large language models: A general framework for the safe implementation of LLMs in clinical care. Interact J Med Res, 14, e59823. https://doi.org/10.2196/59823

**38.** (Zhang, R et al., 2025) Zhang, R., Li, H. W., Qian, X. Y., Jiang, W. B., & Chen, H. X. (2025). On large language models safety, security, and privacy: A survey. Journal of Electronic Science and Technology, 23(1), 100301. https://doi.org/10.1016/j.jnlest.2025.100301

**39.** (Shah et al., 2023) Shah, A., Zafar, A., Wu, J., Shaikh, M. B., Irfan, M., Hadi, M. U., Akhtar, N., Qureshi, R., Mirjalili, S., Muneer, A., & Tashi, Q. A. (2023). Large Language Models: A Comprehensive Survey of its Applications, Challenges, Limitations, and Future Prospects. *[TechRxiv preprint]*. Institute of Electrical and Electronics Engineers. https://doi.org/10.36227/techrxiv.23589741

**40.** (Srinivas et al., 2024) Srinivas, S. S., Vaikunth, V. S., & Runkana, V. (2024). Knowledge Graph Modeling-Driven Large Language Model Operating System (LLM OS) for Task Automation in Process Engineering Problem-Solving. Journal, 222–232.