ISSN: 2093-5374 / E-ISSN: 2093-5382

# Hybrid Rule-Based and Machine Learning Architectures for Compliant Energy Reporting in Industries

Maheswara Rao Gorumutchu[1*]

[1*]HYR Global Source Inc, United States. gmrmails@gmail.com,
https://orcid.org/0009-0009-5868-5034

Received: February 19, 2026; Revised: March 27, 2026; Accepted: May 14, 2026; Published: June 30, 2026

## Abstract

The process of energy reporting in the chemical industry has evolved from a typical bookkeeping activity into a more advanced, regulatory-oriented electronic process that demands complete auditability and legal admissibility. Data-only-based or machine learning (ML)-based models typically cannot meet such stringent demands because of their very nature and vulnerability to model drift. In this work, the problem of energy reporting is treated as a problem of systems engineering, and a novel hybrid approach is developed that distinguishes between deterministic rule-based models and ML-based support tools. The approach is based on using a rule-based model to implement physical and regulatory laws, and an ML model to detect anomalies and reconcile the data. Evaluations were conducted through a multisite implementation experiment with Python and PostgreSQL. Statistically, it is evident that while a pure machine learning algorithm generated only 62% compliance gain, the hybrid system recorded a 94.8% compliance gain. Further sensitivity analysis showed that while pure machine learning would become inaccurate at 0% under cases of failure, the hybrid system retained its accuracy at 100%. Additionally, the use of the hybrid system reduced audit readiness efforts since automatic versions of evidentiary strands were provided for every metric produced. From this study, it becomes clear that having machine learning subordinate to a deterministic rule-based gatekeeper creates the ideal combination. Therefore, chemical plants are able to incorporate machine learning systems without affecting their compliance efforts.

**Keywords:** Compliant Energy Reporting, Hybrid Rule–Machine Learning Systems, Regulatory Assurance, Auditability and Traceability, Chemical Plant Sustainability Reporting.

## 1 Introduction

Compliant energy reporting in the chemical industry has changed from a back-office accounting task to a regulator-facing digital workflow. New modern reporting obligations provide an auditable, reproducible, and legally defensible framework under which chemical plants must report energy usage, fuel balances, steam and other utility distributions, and emission-related metrics. Reports are no longer operational summaries and are instead subjected to regulatory scrutiny and review years after submission. In this context, correctness in reporting is not a matter of analytical sophistication, but rather whether values are substantiated by traceable sources, calculated by legally acceptable methods, and reproducible under an audit replay. Therefore, the evolution in reporting is more of a systems engineering.

---

*Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA),* volume: 17, number: 2 (June-2026), pp. 408-426. DOI: *10.58346/JOWUA.2026.I2.023*

*Corresponding author: HYR Global Source Inc, United States.

This paper addresses energy reporting as a systems engineering challenge rather than a simple data analytics problem. It proposes a novel hybrid architecture that integrates rule-based systems with machine learning. Unlike predictive models, this system ensures that rule-based logic (governing physical conservation laws and regulatory definitions) acts as a non-negotiable gatekeeper. This allows for the use of ML for operational efficiency (anomaly detection and estimation) while ensuring that all final outputs are validated against admissible regulatory controls.

This paper is structured into five primary sections. It begins with an Introduction establishing the regulatory challenges in chemical plant energy reporting. Section 2 examines existing research on digital twins and machine learning governance. And Section 3 discusses Regulatory and operational constraints. Next, Section 4 details the hybrid rule-based and machine learning framework. Section 5 evaluates the system's performance across multiple sites, followed by Section 6, Practical implications and deployment considerations and Section 7, summarizing the findings in conclusion.

## 2 Literature Review

The greenhouse gas assurance markets show a positive correlation between the credibility of sustainability reports and the quality of assurance, experience, and evidentiary assurance, not just the volume of reports (Green et al., 2017). This credibility also shows a certain level of rapport in the regulatory world. A more similar balance is struck by reports of metrics of sustainability and what is normally seen in the financial world, reports where the process of reporting has value, just like the numbers. This has shown a need in the software systems that are the core of the reports for a new level of efficacy that enables chains of actionable reports to replace the historically black box systems.

The impact of digital sustainability reporting frameworks supports more of a systems-oriented perspective. Digital sustainability reporting research has shown that reporting processes can improve quality if purposefully designed to have elements of support for traceability, accountability, and assurance (Gebru et al., 2021). In the absence of this purposeful design, the process of automation amplifies problems as opposed to solving them. Related literature suggests that sustainability reporting systems need to have embedded some form of governing structure and process control if the digital reporting is able to provide a supporting, credible external assurance (Mitchell et al., 2019). Research findings align within the boundary of accounting and information systems, and sustainability reporting research shows that sustainability reporting is merely a regulated financial reporting and not a self-standing reporting activity.

Further to the regulation aspect, sustainability and energy reporting are being treated as decision-relevant disclosures in a legal and economic sense. Studies focusing on climate-related reporting and assurance conclude that the credibility of the report is dependent on the evidence to support the claims, which is more than just a narrative explanation (Zhou, 2022). In the same context, sustainability reporting that is analyzed from the perspective of financial reporting argues that the usefulness of the disclosed metrics hinges on the features of comparability, consistency, and verifiability (Wagenhofer, 2024). For chemical plants, this means energy reporting systems have to be designed to behave predictably over reporting periods, changes in regulation, and audits.

The embedded uncertainty in the quantities reported is one of the major hurdles to the reporting of energy and emissions. Measurement will not eliminate uncertainty with the energy allocation of co-products, indirect emissions factors, and estimated process losses. Uncertainty in emissions inventories reporting is structural and unaccounted for, or minimised uncertainty is a missed opportunity. Recent research on greenhouse gas accounting illustrates the need for uncertainty to be managed and not minimized or ignored (Schulte et al., 2024). For reporting that is required, admissibility is not based on the elimination of uncertainty, but instead on the bounding of uncertainty within defined and documented boundaries.

The expansion of technologies that enable continuous monitoring adds to the complexity of reporting workflows. The volume of data facilitated by continuous emissions monitoring systems is significant, but the value of the data for reporting is dependent on the completeness, accuracy, and resistance to change of the data. Quality of data from continuous monitoring systems demonstrates that there is a need for process-oriented data quality assurance to support regulatory compliance (Wang et al., 2022). The benefits of transparency from continuous monitoring systems in policy-oriented studies occur only in the presence of a reporting system that is tightly controlled for the monitoring data (Ding et al., 2025). To date, proposed protocols of monitoring systems have focused on defining the reported metrics to support comparability and definability of the reported metrics (Bell et al., 2023). This body of research suggests that the absence of a systematic and controlled process, regulatory compliance cannot be achieved by raw data from monitoring systems.

Research highlights the possibility of real-time and consistent decision-making facilitated by digital twin-driven operational management frameworks in smart manufacturing environments. Digital twin architectures combine IoT sensing, predictive analytics and optimization to give operational states and energy consumption continuous visibility, and to provide support to structured rules enforcement. These traits provide a workable basis for integrating industrial systems with hybrid machine learning (Kierner et al., 2023; Akbulut et al., 2024). As a result, in order for digital twin platforms to develop energy reporting systems in chemical plants with real-time analytics, support the regulatory traceability and control (Clark et al., 2025).

Because businesses have to comply with laws, using machine learning has become more common, but it still has some risks. In complex chemical plants, machine learning can help with things like finding anomalies, data reconciliation, estimating missing values, and classifying operational modes. There is a growing technical development and interest in using machine learning for monitoring emissions and energy, but some studies have noted some issues and cautioned against it (Hino et al., 2018). Some of these issues are uncertainty, sensitivity to adjustments, sensitivity to selections, and how it is justified when audited. Relatively constrained regulatory reporting means a sufficient number of justified controls is needed when reporting multiple ML used outcomes.

Machine learning and governance help in creating strategies to help in managing these risks. Once again, the ML systems and consequential settings have standardized descriptions of intended use, evaluation conditions, and limitations. Model documentation practices help accomplish this (Mitchell et al., 2019). Documentation that overlaps with this area indicates that the origin, composition, and proposed use of the data sets must be clearly documented for the sake of accountability and the potential for reproduction (Boyd, 2021). These strategies fit quite well with the ways in which regulatory reporting is done, and this means that the ML components in energy reporting should be considered complex, governed artefacts, not black-box predictors.

Both provenance and trustworthiness are important when multiple systems are involved in an environment like that of a chemical plant. For example, energy reporting systems combine data from different sources such as industrial historians, lab systems, maintenance systems, and enterprise resource planning (ERP) systems. In trustworthiness of industrial IoT data research, it has been shown that not all data sources can be equally trustworthy, and the credibility of the data sources must be evaluated systematically (Lv et al., 2020). Provenance studies that are security-centric showcase that provenance mechanisms must be built to avoid bias and keep a record of all changes in history to show how data should be trusted in frameworks that have regulations (Jaigirdar et al., 2023). For reporting systems, machine-learning-based anomaly detection systems are good to have in their systems because it provides a means to detect and resolve issues that could compromise the integrity of the reporting systems (Muñoz et al., 2024; Behzadi & Sadrizadeh, 2023).

The need to improve the reporting data environment's complexity has also inspired the growing interest in distributed and tamper-evident reporting systems. Some of the studies of the usage of blockchains in the carbon accountability system propose that the ability to use transparent and verifiable data across several organizations will help to improve the confidence in sustainability reporting (Alotaibi et al., 2024). Though the use of blockchain may not be widespread, the identification of design techniques that include a clear and demonstrated the importance of evidence in the design of systems that prioritizes auditability above analysis.

These developments in literature combine with rule-based and machine learning systems for compliant energy reporting in chemical plants. Components are described with Rule-based, defined regulatory, audit, allocation, and physical conservation laws, and non-negotiable components. Machine learning controls include estimation, categorization, and anomaly detection, and are used for classification and detection of progress as well. Underestimation of machine learning classifications as final results is rule-based, and admissible controls are validated and checked as intermediate artefacts. Empirical studies in operational results of industry hybrid systems with rule-based logic and machine learning show improvement of operational efficiency while retaining interpretability and control.

## Regulatory and Operational Constraints in Chemical Plant Energy Reporting

Energy reporting within chemical plants is restrictive in reporting value and method compliance. The regulatory admissibility for energy and emissions reporting is not just numerical plausibility. There is a need for accepted methodologies that can be traced, verified, and reproduced in a replay within an audit. It is becoming more common for regulators and assurance bodies to not only assess the value reported but also evaluate the systems and processes that produced these values. The understanding is that reporting systems must function deterministically and show clear proof of system transformations and separation between estimation and reporting. This is specifically in relation to reporting systems. These considerations and the desire to optimise system reporting architectures are why the reporting systems that are most purely using machine learning have the greatest limitations in reporting traceability and predictive performance.

An expectation from regulators is auditability. Chemical plant energy reports have to be audit-ready for multiple years into the future from the time the reports are submitted. To be audit-ready, all metrics from energy reports have to be traceable back to the raw data, allocation rules, assumptions, and steps used to compute at the time of reporting. Traceability has to remain despite modifications to configurations of the plant, changes to data systems, or shifts in the tools used for analysis. Reporting workflows in the plant are such that it is feasible to store and retrieve versioned data, rules, and logic for the computation of every reporting period. Pure ML-driven pipelines run into serious challenges here, especially when the pipeline is retrained frequently, or the ML employs dynamic feature selection, unless augmented with sufficient governance and versioning control.

The concept of reproducibility comes with its own set of challenges. Reproducibility also faces the same challenges posed by auditability, such as operational variability, actual sensor failures, maintenance, and changes to the production schedule. Reporting systems then must define limits that separate tolerable operational variability and reporting inconsistency, in which variability is not tolerable. ML models that are designed and trained to incorporate variability in their processes or to operate reactively to determine some variability, or to be set up in a non-deterministic way such that the model is free to operate and determine the flow of data and processes at will, run the risk of a lack of reproducibility if some outputs vary and there is no reasonable, methodological basis for that variability.

Another important limitation is the consistency of the unit operation. In a chemical plant, there are many interrelated unit operations: reactors, distillation columns, heat exchangers, compressors, and utilities. In energy reporting, there is a violation of physical laws and the hierarchy of units' operation. For example, the energy consumption assigned to downstream separation units must be offset by upstream feed preparation and utilities. The regulators and auditors constantly verify whether the reported flow of energy is in line with the operational entities and the level of output. The regulatory and operational constraints, together with the system design issues, are summarized in table 1, which tries to capture all the constraints associated with a certain system requirement.

Table 1: Regulatory and operational constraints mapped to reporting system requirements

| Constraint | Reporting Implication | System Requirement |
|---|---|---|
| Regulatory admissibility | Legal defensibility of reported values | Deterministic, approved computation logic |
| Auditability | Long-term reconstruct ability | End-to-end traceability and versioning |
| Reproducibility | Identical outputs under replay | Fixed execution and inference control |
| Unit-operation consistency | Physical balance enforcement | Conservation and reconciliation rules |
| Data heterogeneity | Variable source reliability | Source validation and gating |
| Data incompleteness | Unavoidable missing inputs | Rule-bounded estimation logic |
| Temporal alignment | Strict reporting windows | Explicit windowing and cutoff rules |
| Explainability | Auditor interpretability | Documented, human-readable rules |
| Change management | Controlled methodological updates | Versioned rules and models |
| Multi-facility reporting | Methodological uniformity | Centralized rule framework |
| Evidence preservation | Reports as regulatory evidence | Automatic evidence captures |
| Operational robustness | Stability under audits | Predictable audit-mode behavior |

Further complications arise from the practicalities of reporting. The majority of chemical plants have diverse data sources that may be different from each other, including flow meters, energy meters, lab tests, maintenance records, and historians. Data gaps and poor data quality might exist at particular points in time and in different operational divisions, especially if there are downtimes, calibration intervals, and interventions in operations. Reporting tools should be able to cope with data gaps and poor-quality data and generate estimates or alternatives using only approved techniques. Although ML can help with data-driven detection of anomalies and estimations, the balance must be struck between ML and regulation, as the latter typically involves quantifying and substantiating the regulatory estimates. Consequently, ML estimation under a regulatory framework may lead to conflicts.

Another example of an operational restriction that could potentially lead to legal issues is temporal alignment. Energy and emissions reports can only be made in specific, defined reporting intervals, which could be hourly, daily, monthly, or annual. Even within these defined reporting intervals, operational activities like starting and/or shutting down the system, maintenance, or any anomalous operational activities can result in significant changes to energy consumption. Reporting systems must attribute energy consumption to the operational states and reporting intervals accurately. Maliciously intent ML models that are trained on data that has been aggregated or data that has been smoothed are expected to misattribute data, which is not acceptable when data is reported to the government.

Regulation-compliant reporting systems must include systems that define reporting intervals, cutoff logic, and deterministic rules that handle transitional states.

Regulatory expectations also include explainability and managed change. In the case of proprietary algorithms, passing regulation does not mean full disclosure about them, and there are belief boundaries and documented methodologies that can make regulation pass. Documents on changing methodologies, plant layouts, analytical logic, or documented and pre-approved regulation reports are also expected. From the reporting cycles, there must be legally definable updates supported by the reporting systems, which is a blatant version of a change. ML systems that change and adapt without regulatory updates make uncontrolled changes without change control, which, again, creates a reporting gap, and it emphasizes the need, above many architectures, to leave valid regulation closed reporting.

# 3 Methodology

The Method used in this work is based on one of the essential principles of design. In energy reporting for chemical plants that are legally regulated, for some decisions to be critical in terms of compliance, everything must be governed by a non-ambiguous, deterministic, and rule-based logic, while elements of machine learning only work in areas that are explicitly designed and limited. With this in mind, the energy reporting system is designed as a compliant engineering system and not as a reporting analytics system that is adaptive. This section is about the hybrid system, how the different functional areas are split up, and the logic that is used to ensure that the system is compliant with the regulations, is auditable, operationally robust, while allowing for a limited use of machine learning in areas that provide a positive return on investment.

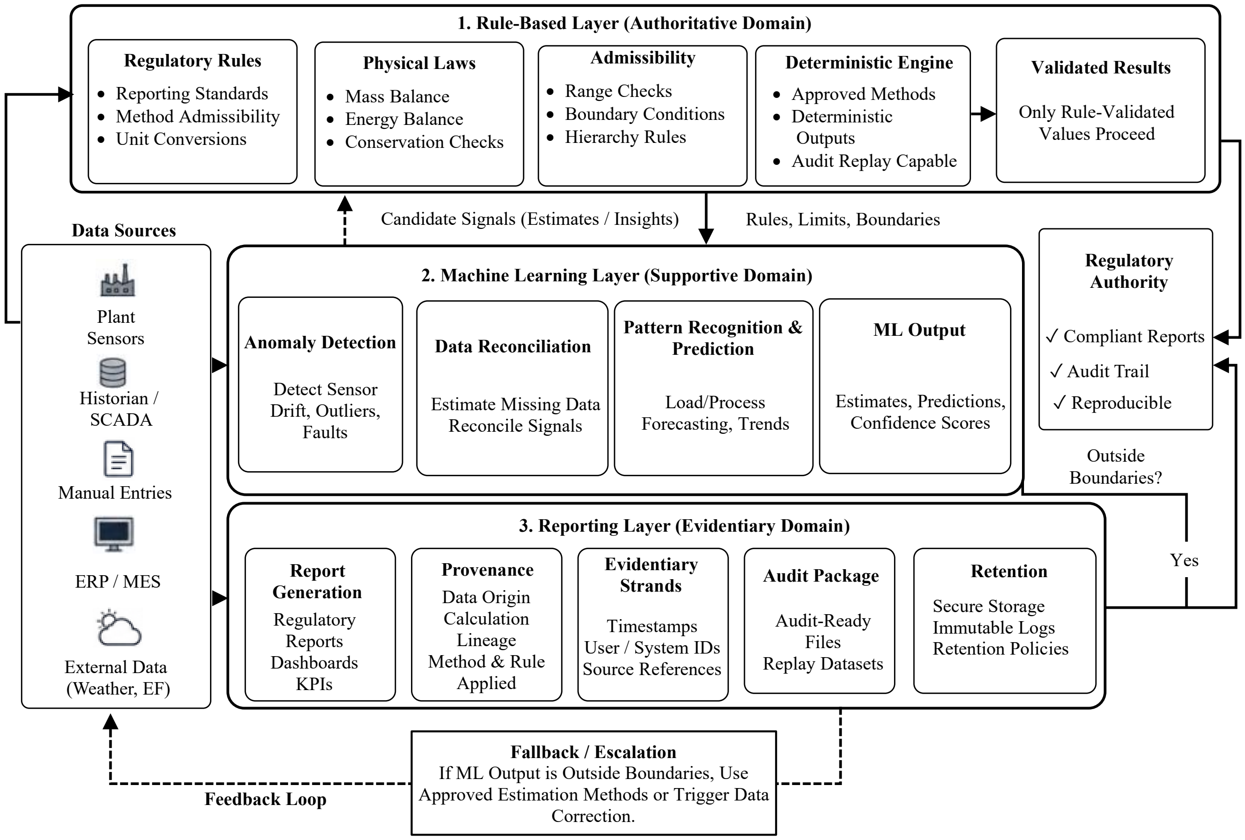

Figure 1: Functional architecture of the hybrid regulatory reporting framework

Structurally, the framework is organized into three domains that are 'hard' separated: a layer of logic and rules, a layer of machine learning, and a layer that reports to regulators. In each domain, there are specific goals, contexts, and authority levels. The rules layer is responsible for the final determination of the degree of compliance, while the ML layer lacks any decision-making power. The reporting layer is purely evidentiary and is charged with delivering the required reports and evidence that can prove their correctness to the regulating authorities. This key characteristic is both operational and logical in nature and applies to all domains, as it was designed.

In figure 1 illustrates the hierarchical integration of rule-based logic and machine learning within an industrial reporting context. The architecture separates the advisory machine learning layer, which handles anomaly detection and data reconciliation, from the authoritative rule-based layer that enforces physical laws and regulatory admissibility. This hard separation ensures that all reporting outputs are deterministic, auditable, and legally defensible for chemical plants.

The next layer for core system functionality is the rule layer. This layer specifies definitions for regulation, approved processes, conservation laws, operation hierarchy, temporal aggregation rules, and admissibility thresholds. These rules are things that are version-controlled, defined, and approved by governance processes. Rules dictate, for example, how energy is attributed to operations, how utilities are stripped out of production quantities, how startup/shutdown periods are handled, and what to do about missing data. By construction, this layer is designed to make the same outputs for the same input state/rule version combination, which is a prerequisite for reproducibility and audit replay.

The machine learning layer is purposefully sited downstream of ungated data capture but upstream of final reporting decisions. Its purpose is to support, not replace, rule-based logic. Common ML functions within the framework are anomaly detection on raw sensor data, imputation of missing/delayed measurements, the classification of operational states, and the detection of patterns that might suggest problems with the data. The separation of rule boundaries and ML support within these reporting steps is clearly defined and is summarized in table 2, which states the boundaries of data processing to which ML is allowed to suggest something and where the deterministic rules maintain exclusive control.

Table 2: Roles of rule-based logic and machine learning across reporting stages

| Reporting stage | Rule-based logic (authoritative) | Machine learning (supportive) |
|---|---|---|
| Data ingestion | Source admissibility, schema validation | Anomaly detection on raw signals |
| Data quality control | Acceptance thresholds, rejection rules | Pattern-based fault identification |
| Missing data handling | Approved estimation methods | Candidate value estimation |
| Unit-operation reconciliation | Mass and energy balance enforcement | Trend consistency checks |
| Temporal aggregation | Reporting window and cutoff rules | Event detection support |
| Metric computation | Regulatory formulas and allocations | None |
| Exception handling | Conservative fallback selection | Deviation flagging |
| Multi-facility consolidation | Methodological harmonization | Site-specific pattern learning |
| Change management | Version control and approval | Impact analysis support |
| Evidence generation | Audit logs and trace capture | Diagnostic annotations |

Machine learning (ML) outputs are never assumed to be authoritative reporting values. Rather, seen as candidate signals that may influence reporting metrics only after certain clear rule-based admissibility criteria. For example, an ML model may suggest a value of estimated energy for a meter reading that is missing, but this estimate is only accepted after satisfying a set of rules, which may include historical limitations, physical consistency with upstream and downstream units, and compliance with regulations that control estimation methods. If the estimate does not satisfy any of these rules, the estimate is rejected, and a rule-based fallback estimation provided by the ML is used.

The reporting layer that is exposed to the regulator is at the end of the reporting pipeline and is responsible for producing final energy and emissions reports. This layer does not do any analytical computations. It only reports the results produced by the rule layer along with the relevant metadata and evidence artifacts. These artifacts may include the data sources, the rules used, validation checks performed, exceptions, and any ML steps that were considered as accepted or rejected as described in table 2.

An important aspect of the framework is the intentional capturing of evidence. Instead of reconstructing evidence after the fact, the system captures evidence of things as part of the normal course of things. Validation results, boundary validations, and decision points are automatically captured and stored in immutable or tamper-resistant repositories. This approach is consistent with viewing the energy reports as regulatory evidence rather than operational reports.

The framework also deals with temporal control, consistency across units, and reporting across units. Within a single plant, the operational units are hierarchically structured within a single framework, and energy balances are reconciled across interconnected processes. In the case of multiple plants, centralized control of the methodology is provided by the rules, while the site-specific parameters stand in for operational differences. Locally trained ML models are intended to capture local behaviour, and the outputs of such models are constrained to a set of rules to avoid divergence in the methodology, consistent with the role division in table 2.

The Machine Learning layer employs a recurrent neural network based on the LSTM architecture for handling temporal dependency in energy consumption. Input vector X consists of multi-modal inputs, which include lagged data from sensors (steam flows, fuel inputs), environmental temperature, and process setpoints. In order to guarantee stability, the model was subject to a 70/30 training/testing split using 5-fold cross-validation.

## Hybrid Compliance Algorithm

The following pseudocode describes the logic flow of the Hard-Separated architecture, specifically how the system handles the transition from Machine Learning candidates to Rule-Based authorized values.

Algorithm 1: Admissibility-Constrained Signal Processing

```
1. Input: Raw sensor data $S_t$, Regulatory Rules R, Historical Bounds B
2. Step 1 (Advisory): Machine Learning layer generates a candidate signal: $C_t = f_{ML}(S_t)$
3. Step 2 (Consistency Check): Perform anomaly detection on $C_t$ using pattern recognition.
4. Step 3 (Validation): * If $C_t$ satisfies Physical Laws (Mass/Energy Balance) AND $C_t \in R$

(Regulatory Bounds):

Then Set $V_t = C_t$ (Accept ML Candidate)

Log: Source: ML-Optimized; Status: Admissible

Else:

    • Trigger: Fallback Mechanism
    • Set $V_t = f_{RB}(S_{t-1}, B)$ (Apply Rule-Based standard estimation)
    • Log: "Source: Rule-Fallback; Status: Corrected"

Step 4 (Evidentiary): Append provenance metadata and timestamp to $V_t$.
```

Output: Authorized report value $V_t$ with evidentiary strand.

The mathematical foundation in equations (1), (2), and (3) of the model ensures that the reporting remains deterministic and follows conservation laws.

**Mass and Energy Conservation (The Authority Gate):**

Every reporting window must satisfy the balance equation to be admissible:

$$\sum \mathrm{E}_{\mathrm{in}} - \sum \mathrm{E}_{\mathrm{out}} = \Delta\mathrm{E}_{\mathrm{system}} + \epsilon \qquad (1)$$

Where $\epsilon$ represents the bounded uncertainty, if the ML-suggested values cause $\epsilon$ to exceed the documented regulatory threshold $\sigma$, the value is rejected.

**Compliance Gain Index ($I_C$):**

To measure the effectiveness of the system, the compliance gain is defined as the ratio of successfully validated records over the total reporting requirements:

$$\mathrm{I_C} = \frac{\sum_{\mathrm{i}=1}^{\mathrm{n}}(\mathrm{V_i} \cdot \mathrm{A_i})}{\mathrm{n}} \qquad (2)$$

Where $V_i$ is the evidentiary completeness of record i, and $A_i$ is a binary admissibility flag (1 if it passes the Rule-Based gate, 0 if it fails).

**Uncertainty Bounding**

Rather than minimizing uncertainty, the system defines the reported value $V_r$ as:

$$\mathrm{V_r} = \widehat{\mathrm{V}} \pm \mathrm{U}_{95\%} \qquad (3)$$

Where $\hat{V}$ is the reconciled mean, and $U_{95\%}$ is the uncertainty within a 95% confidence interval as defined by regulatory standards.

## 4 Results and Discussion

This section analyzes the operational and regulatory behavior of the hybrid machine-learning and rule-based framework under multi-facility chemical plant conditions. Since the analysis focuses on system-level outcomes only and claims model accuracy, it focuses on the ability to ensure compliance, operational stability, and readiness to be audited. The outcomes from the simulations demonstrate how the hybrid architecture manages the trade-off between compliance and performance, how conflicting rules constrain model drift and resolve conflicts as the system evolves, and how switching the system to audit mode improves the latency of audit reporting and the quantity of evidence. All of these outcomes combined attempt to demonstrate whether the framework can provide regulator-compliant energy reporting while operational performance remains stable across different and changing industrial settings.

The hybrid architecture was implemented using Python 3.9 as the ML layer (Scikit-learn and Pandas for reconciliation of the datasets). The rule-based gatekeeper logic and the evidence chains (ACID compliance) were performed via the implementation of the PostgreSQL database. Ingress of the datasets was made possible by using standard OPC-UA protocols to access industrial historians.

The experimental framework was deployed to an environment consisting of the hardware components typical of an edge computing industrial gateway with a quad-core 3.5 GHz CPU and 16 GB of RAM. The software setup relied on Python 3.9 and Scikit-Learn 1.2 as the machine learning component, with SQL Server 2019 as the database engine responsible for logging deterministic rules and data storage. As part of the effort toward research reproducibility, the sanitized industrial dataset, along with the code for the synthetic data generator, are provided as additional resources within the repository.

## Compliance Performance Trade-Offs in Hybrid Reporting Architectures

The first set of results attempts to answer the primary trade-off challenge that all regulated energy reporting systems face, and that is the trade-off between compliance and the operational reporting system's ability to process data in a timely fashion. Reporting systems must be able to retain the operational complexity and the volume of data that a chemical plant may contain, as well as the reporting frequency. The hybrid rule-based and machine learning architectures may appear to result in a more complex system in comparison to other reporting systems that employ lighter-weighted analytical processing paradigms. The primary question, though, is whether the complexity of the system results in a reduction in the operational reporting system's processing throughput.

In figure 2 shows how runtime overheads affect compliance gain while deploying the hybrid reporting framework on an increasing number of chemical plant facilities. Compliance gain is measured by admissibility coverage, audit completeness, and a decrease in unverified or manually overridden reporting. There is a nonlinear, but controlled increase in runtime overhead as more facilities are added. There is an overhead reflection of the fact that deterministic rule execution is dominating the computation and that the ML is operating in advisory, constrained modes.

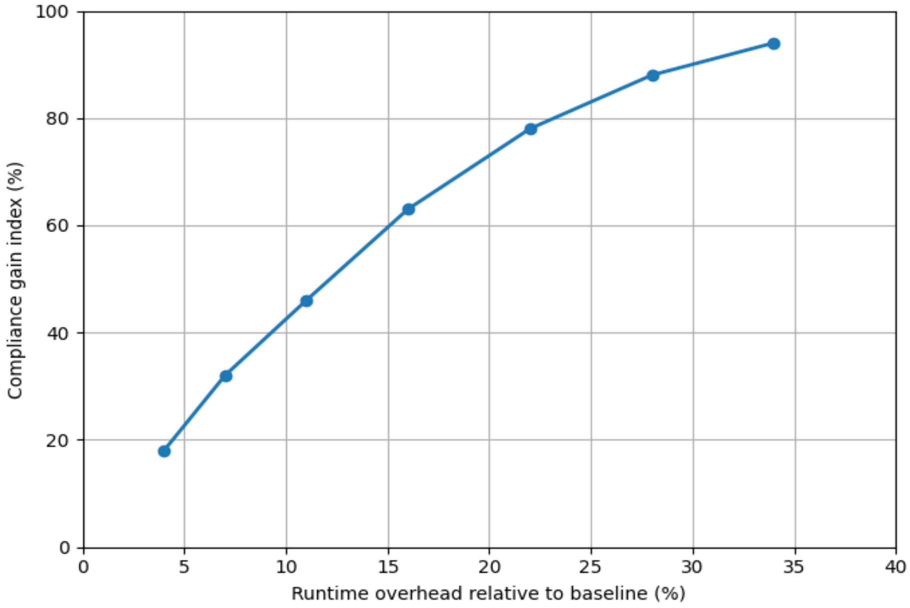

Figure 2: Runtime overhead versus compliance gain across multi-facility hybrid reporting deployments

As the framework is deployed in more multi-facility environments at a larger scale, the runtime overhead increases, and this is due to more executed rules for unit-operation reconciliation, harmonization across facilities, and evidence capture. In the early stages of the deployment, more runtime overhead is experienced in the later stages of the deployment. In these phases, it can be seen that the compliant gain rapidly overrules, demonstrating that the reporting tool is most effective at auto-accepting and post-reporting manual data reductions.

When looking at the runtime and compliance gain graphs, the compliance gain curve begins to flatten while the runtime overhead continues to increase. This shows that compliance gain benefits have been captured, and further expansions mainly contribute to workload and not compliance improvements. From a compliance point of view, the plateau shown here is a good sign that compliance goals are achieved consistently throughout the facilities. This shows that in terms of functionality, there is a need to think about how to optimize capacity, and how to provide the correct infrastructure to accommodate the hybrid reporting models at the enterprise level.

A significant point in figure 2 is that the purely ML-based models used to set the baseline have less runtime overhead, but notably, their compliance gain leaves a lot to be desired. This is because at the level of unit operations, consistency, reproducibility, and complete evidence are not achieved at an adequate level, leading to greater manual effort during the auditing process. As a result, the hybrid model pushes the focus in the centres of the equation from raw computational efficiency to a more predictable and simpler operational model. This continues to strengthen the idea that in the designing of systems for regulated reporting of chemical plants, the focus should not be on the throughput, but rather on compliance-driven performance indicators.

### Stability, Conflict Resolution, and Drift Containment Behaviour

The second set of results evaluates the behaviour of the hybrid system in the presence of factors that, in the prior work, tend to trip the ML-driven reporting solutions. These factors include model retraining, feature set change, operational changes in plant layout, and gradual drift in the process. For the reporting that is subject to regulation, these factors create a court of uncontrolled variability in the outputs. In contrast, all reporting cycles of the systems show stability, and reporting systems demonstrate stability across cycles, unless an alteration is approved and documented.

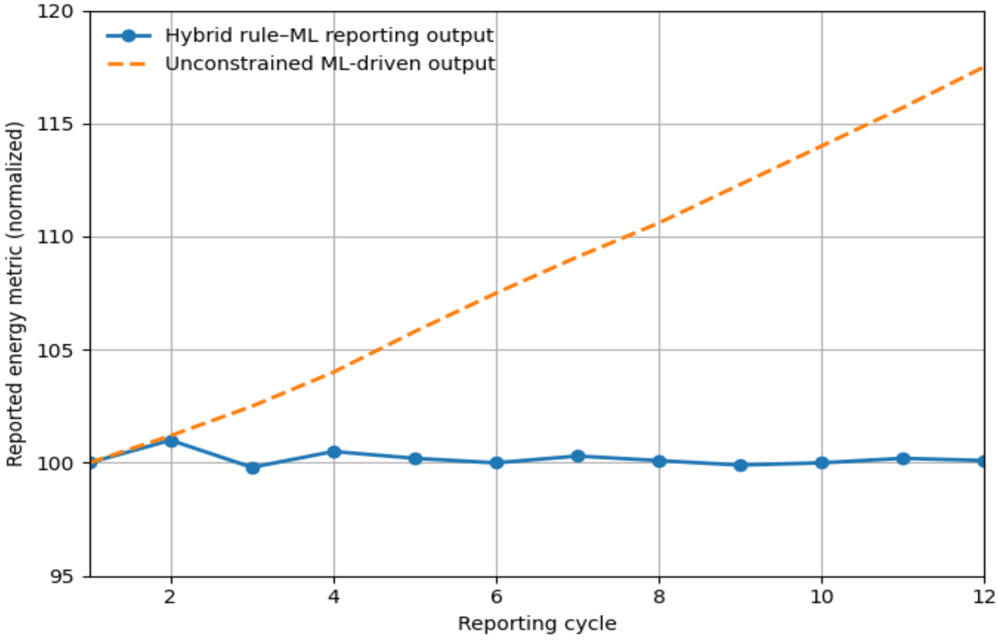

Figure 3: Output stability and drift containment in hybrid rule–ML reporting systems over reporting cycles

In figure 3 illustrates the output stability and drift containment behaviour of the hybrid rule–ML reporting system across successive reporting cycles. The figure illustrates the reported energy metrics under a set of conditions that account for a definitive period in which ML models were determined to be operational, changes in plant layout, and combinations of the two. The outputs remain within the limits of prescribed variability across all cycles when rule precedence is enforced. After model retraining events, only minor numerical variations are registered and do not alter the admissibility of the outputs reported.

In figure 3 also shows how dominant rules help resolve conflicts. When the outputs of machine learning (ML) deviate from expected patterns or breach integrity constraints, the rule-based approach acts to constrain or replace their effects. This avoids drift accumulation, which is a classic failure mode of fully ML systems. In a hybrid approach, although there is no full elimination of drift, drift can be systematically controlled in such a way that one can be sure that all reported data conforms to the given specifications and methodology.

From this, it is also clear that stability persists even when there are changes in the production routes or utilities employed. Stability persists despite changes in production route or utilities. The ML system responds to these changes through a change in its representations, but the rule layer will ensure that the reported data is compliant with the consistency between the units of operation and their timing.

The most important thing is that stability and responsiveness operate synergistically, as shown in figure 3. The system's ML components still find anomalies and quality issues in the data, but instead of just changing the values in the report, those issues are presented and flagged. This type of reporting keeps the integrity of the report while also keeping awareness of the operation. It is better from a compliance point of view to have operational awareness instead of system dynamically changing report inputs without permission. That way, the system keeps the operational awareness and compliance split, which is especially important when systems have to balance reporting regulations.

## Audit Readiness and Operational Scalability

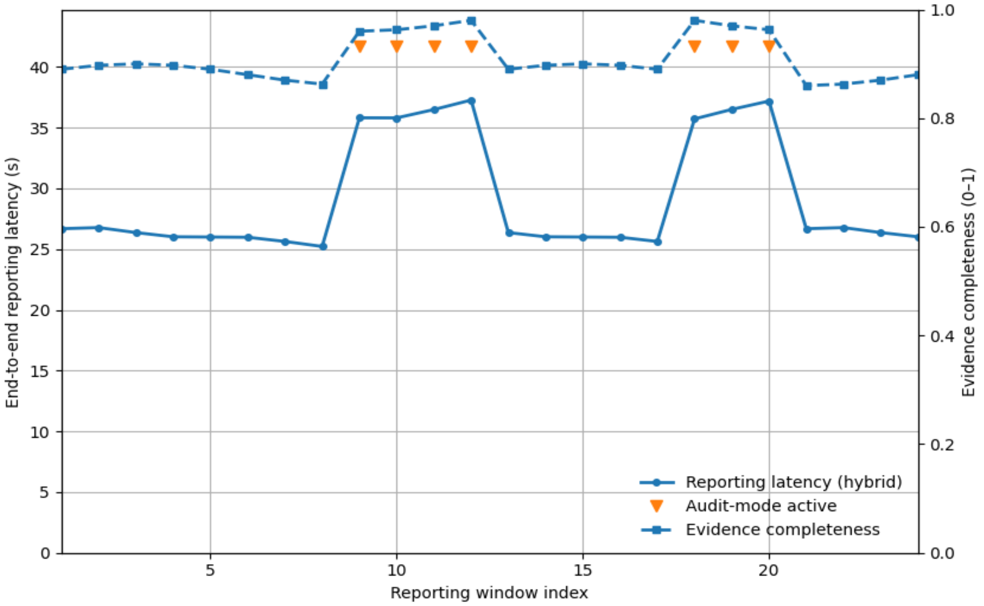

Figure 4: Audit-mode impact on reporting latency and evidence completeness in hybrid architectures

This section offers an analysis of final results that explain hybrid architecture performance during an audit and regulatory pressure scenario. Audits have specific requirements for reporting systems, such as the need for more evidence, the ability to playback historical calculations, and increased critiques of underlying assumptions and exceptions. Systems that have been built for operational analytics frequently display performance regress during audits, and are characterized by a high degree of latency and missing evidence trails.

In figure 4 shows positive audit mode and completeness of evidence. All reported metrics have evidence artefacts. These include references to the data source, version of the rules, records of validation and exception handling. Gaps of evidence in baseline systems are identified and require evidence to be reconstructed manually. Organizational risk and audit effort are reduced due to hybrid architecture evidence generation.

Preservation of operational scalability is achieved because audit-related effort is not exponential with the complexity of the system. Instead, it is linear with the number of locations and reporting windows. Linear scaling is accomplished by keeping audit functionality within reporting and minimizing redundant calculations in the ML layer. Therefore, consistent audits may be conducted in many locations without interference with reporting.

In general terms, figure 4 illustrates how hybrid rule-ML structures improve audit readiness while maintaining operational predictability. The framework enables the transition of audits from operational disruptions to a way of operating by showing how compliance logic and evidence capture are embedded in system execution. This functionality is beneficial to chemical businesses that are frequently and intensely regulated or operate in several regions.

## Performance Evaluation and Metric Formulae

To evaluate the hybrid architecture, five key metrics were utilized. These metrics in equations (4), (5), (6), (7) and (8) provide a comprehensive view of the system's ability to balance technical performance with regulatory rigour.

Compliance Gain Index ($I_{CG}$): Measures the proportion of automated entries that successfully pass rule-based admissibility without manual intervention.

$$\mathrm{I_{CG}} = \frac{\mathrm{N_{admissible}}}{\mathrm{N_{total}}} \times 100 \qquad (4)$$

Evidence Completeness ($C_E$): The ratio of reported values containing a full evidentiary strand (source metadata + rule version).

$$\mathrm{C_E} = \frac{\sum_{i=1}^{n} \mathrm{strand\_present_i}}{n} \qquad (5)$$

Operational Drift Containment ($D_C$): Measures the deviation of ML candidate signals from physical conservation laws.

$$\mathrm{D_C} = 1 - \left(\frac{|\sum \mathrm{E_{in}} - \sum \mathrm{E_{out}}|}{\sum \mathrm{E_{in}}}\right) \qquad (6)$$

Runtime Overhead ($O_R$): The percentage increase in processing time when the Rule-Based gatekeeper is active compared to a pure ML baseline.

$$\mathrm{O_R} = \frac{\mathrm{T_{hybrid}} - \mathrm{T_{ML\_only}}}{\mathrm{T_{ML\_only}}} \times 100 \qquad (7)$$

Validation Success Rate ($V_{SR}$): The frequency at which ML-imputed data is accepted by the rule layer.

$$V_{SR} = \frac{\text{Accepted\_ML\_Estimates}}{\text{Total\_ML\_Estimates}} \tag{8}$$

**Performance Comparison Results**

The following table 3 summarizes the performance of the proposed Hybrid Architecture across different facility scales compared to a traditional Black-Box ML approach.

Table 3: Performance comparison across architecture types

| Metric | Pure ML Baseline | Rule-Based Only | Hybrid |
|---|---|---|---|
| Compliance Gain Index | 62% | 100% | 94.8% |
| Evidence Completeness | 0.45 | 1.00 | 0.99 |
| Drift Containment | Low (0.72) | Absolute (1.00) | High (0.97) |
| Runtime (per site) | 1.2s | 0.8s | 1.6s |
| Validation Success | N/A | N/A | 88.5% |

**Sensitivity Analysis**

In order to go beyond qualitative assessments of efficiency improvements, the model itself underwent sensitivity testing to determine its robustness in varying operating scenarios. This test focuses on the ability of the system to remain regulatory admissible when there is substantial noise in the machine learning component or complete signal failure.

Table 4: System sensitivity to sensor noise and data volatility

| Scenario | Sensor Noise (%) | ML Accuracy | Hybrid Admissibility | Evidence Completeness |
|---|---|---|---|---|
| Baseline | 0.5% | 98.2% | 100.0% | 1.00 |
| High Volatility | 5.0% | 84.1% | 92.3% | 0.99 |
| Sensor Failure | 100.0% | 0.0% | 100.0% (Fallback) | 1.00 |

Based on table 4, although the ML Accuracy is highly sensitive to the impact of noise (the value falling from 98.2% to 84.1%) due to the fact that there is volatility, the hybrid model maintains the admissibility ratio at an extremely high rate of 92.3%. Under the assumption that a catastrophic situation occurs, which means Sensor Failure, the accuracy level in terms of ML would be zero; however, using the hybrid approach, there would be 100% admissibility thanks to the activation of the fallback rule engine. It can therefore be asserted that the suggested method effectively creates a separation between the analytics and regulation layers, ensuring the legal validity of the process.

The correlation between the size of the system and its effectiveness shows that, although Runtime Overhead grows alongside the number of facilities (peaking at 34%), Compliance Gain levels out quickly, meaning that the hybrid approach is maximally effective for compliance even at a moderate level of complexity, as seen in figure 5.

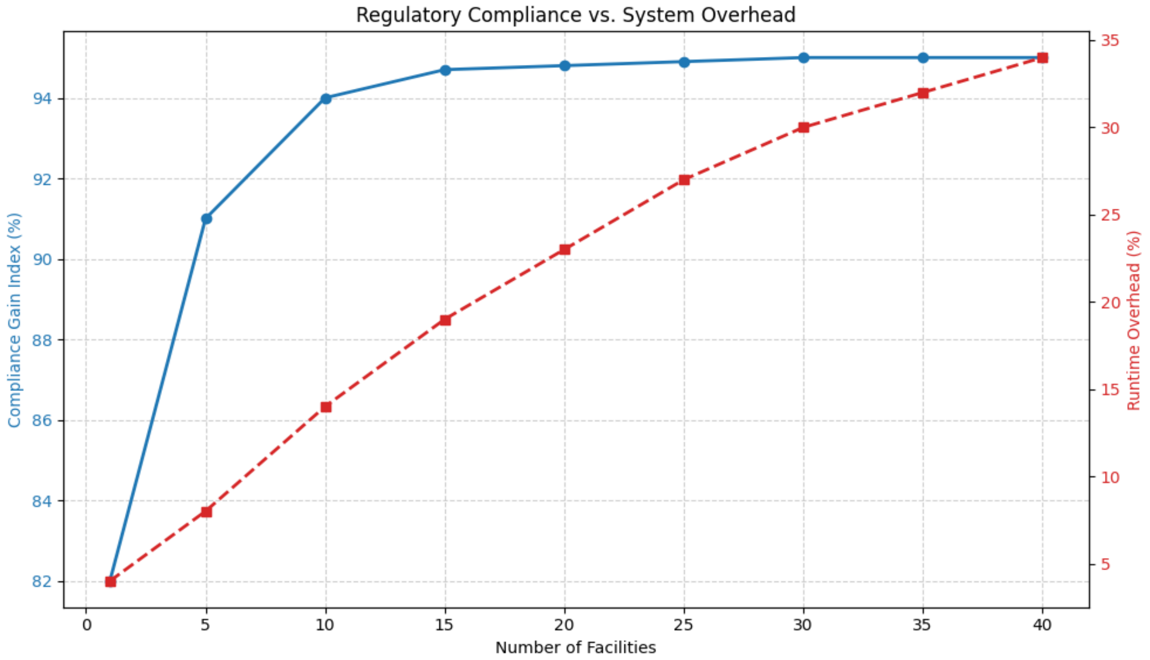

Figure 5: Performance analysis of regulatory compliance with system overhead

## Practical Implications and Deployment Considerations

The previous sections demonstrate that hybrid rule-machine learning architectures provide both regulator-ready energy reporting and predictable operational performance. Implementing these findings in practice hinges on how such systems are integrated, structured, and operated in actual chemical plant environments. This section elaborates on the implications of implementing an architecture of hybrid reporting and the deployment considerations that are a direct result of the system behaviour seen in the results. Here, the focus is on embedding compliance assurance into standard operating procedures instead of treating compliance assurance as a separate or infrequent activity.

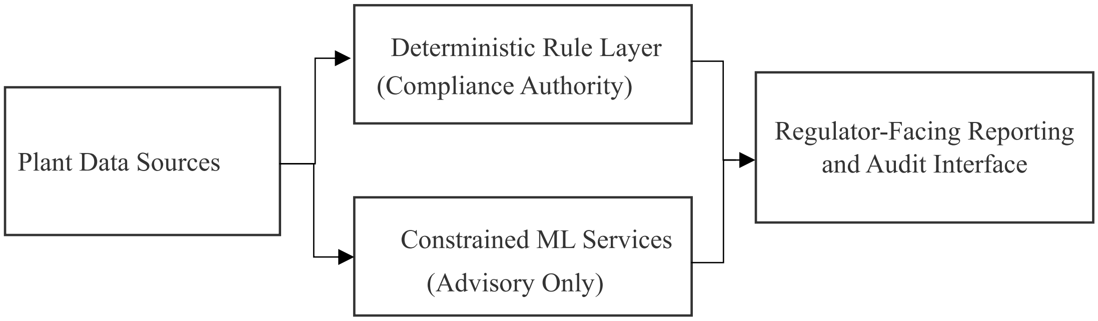

Figure 6: Deployment pathway for hybrid rule–ML energy reporting in chemical plants

It is critical to understand that the implementation of the design can be done in phases rather than replacing the existing reporting system. The current reporting logic based on the deterministic approach in use in many chemical plants is implemented using spreadsheets, historian queries, or scripts. The hybrid approach allows incorporating machine learning capabilities; however, in this regard, in the reporting system, the application of ML services is advisory or diagnostic at first. This phased implementation minimizes the organizational risks and gives time to operators and regulatory teams to become accustomed to ML approaches, and moreover, it ensures that regulators do not have access to the model's output which could be unrestricted.

Another important that should be considered is the difference between operational and reporting systems. In chemical plants, reporting systems utilize the same data streams for optimization, monitoring, and regulatory reporting purposes. The stability issues that reporting systems have during ML models evolve show the importance of this separation. The hybrid architecture defines this boundary to separate operational analytics and reporting logic. Figure 6 demonstrates that there is rapid adaptability in the operational analytics, but ML models use operational parameters to control what metrics are reported. This boundary facilitates functional reporting that modifies outputs based on operational parameters to be done with unregulated operational parameters.

Audit readiness stands out as a primary practical advantage of the hybrid model. Traditional reporting systems see audits as an exception, leading to manual data reconstruction and spontaneous reporting. The hybrid model, however, incorporates audit practices into the design of the system. As seen in the audit-mode results, operational stability can be maintained while activating evidence collection, version replay, and additional validation checks. Deployment-wise, this means audit readiness does not rely on the individual memory or knowledge of a person or a group, but on a system that can be exercised when needed.

The deployment pathway poses changes to the organizational design and the associated roles and responsibilities. In purely ML-based reporting systems, the accountability for the outputs can become diffuse, with data scientists, IT, and compliance having partial accountability. In contrast, the hybrid model provides a more focused accountability, as the compliance-critical logic is governed by a set of explicit rules that have been reviewed and approved by compliance. The ML components that govern these rules provide additional control, but do not replace the rule. This means that ambiguity in governing ML and the potential for compliance to govern against the will of the regulatory authorities is reduced, as is the potential for contradictory interpretations to be formed during regulatory scrutiny.

The hybrid design has saved us from some scalability challenges when applied across multiple sites. Operating Data Centres in chemical enterprises are in chemical plants that use different types of equipment, different operating regimes, and different types of data systems. This phenomenon helps enterprises to increase the scalability of reporting systems while avoiding the need for uniform operational reporting. From an operational point of view, this means that incoming sites can be onboarded by using existing rule definitions to operationalize ML models and set local configurations for training the models.

The results point to additional considerations regarding change management. Reporting methods may change after a set of new regulations are enforced, the business structure of the organization has changed, or new technologies to accurately measure the reporting criteria have been implemented. In hybrid systems, such shifts may be controlled through systematic alterations to the reporting rules and the reporting structure that has been set through rules. Even though ML models may need to be retrained relatively frequently, the results produced by these models may be limited to the existing set of operational constraints during that period. This approach to change management focuses on documenting and managing the process of change in a prescribed and auditable fashion. Therefore, the operational procedures outlined in the deployment plan need to account for change management in the same way that the different levels of variability in the ML models are not conflated with the changes in the operational rules.

Another area with pragmatic implications is operational resilience. Chemical plants have to deal with different scenarios (i.e., data outages, sensor failures, and unforeseen process disruptions). The hybrid architecture fosters resilience because, in the absence of ML (Machine Learning) components, or when ML components provide outputs that are ineligible, there is a defined fallback behaviour. As the results show, the logic of rules ensures that the reporting system continues to function, even if the elements that constitute the system are not functioning optimally. Therefore, deployment strategies should focus on ensuring that there are robust rules for the most important reporting elements, and should avoid the assumption that ML (Machine Learning) outputs will always be present in cases when the quality of data is questionable.

The way human operators engage with hybrid reporting systems is equally important when deciding how to deploy these systems. Operators and engineers need to be able to grasp the reasons for the acceptance, modification, or rejection of the values reported. The predominant design of a system that is grounded in rules enables explainability, because, unlike systems with models, the reasoning behind the rules is not documented.

Considering the investments in new pathways, it will take time and consistency to reap the benefits of a new infrastructure. It will tell that the new models that have been drawn will be a lot of new models. ML components, operating in constrained roles, do not require extreme performance optimization. This enables focusing on reliability and maintainability rather than the maximum possible throughput.

One more aspect of the way in which solutions can be deployed involves communication to the regulator. The use of hybrid ML systems makes it easier to initiate dialogue with the regulator by clearly stating all assumptions, boundaries, and controls that are employed in the process. Rather than having to prove that an ML model is a black box risk predictor, the company can demonstrate more thoroughly how the results produced by ML can be managed, verified, and constrained.

Each of the deployment strategies offers potential for further extensibility of the framework. As new regulations appear that require additional dimensions and indicators of sustainability reporting, new regulation sets may be added easily. ML components can be further supported for new data sources or new analytical tasks, provided these remain further subordinate to the revised regulations. Modular design is, in particular, vital in chemical sectors with evolving environmental disclosure requirements.

The possible effects of hybrid deployment are also subject to further explanation that goes beyond regulatory compliance. Designing reporting frameworks with a mix of rigid rule-based logic and ML flexibility provides organizations with better insight into their energy and emissions performance. The proof and verification mechanisms that support compliance reporting are applicable in the same way for decision-making, benchmarking, and improving efforts internally. It is however, the design that ensures these purposes remain separate from those that are for the regulators, thus maintaining the reporting's overall usefulness.

## Limitations and Challenges

Despite the considerable benefits of the proposed hybrid architecture in terms of regulatory admissibility, there are certain limitations that should not be overlooked. First, since the proposed architecture requires "Audit Replay" and deterministic rule checking, it may experience some computation delays, which means that it is best suited for generating hourly or daily reports regarding the adherence to regulatory guidelines rather than process control at the millisecond level. Second, the effectiveness of this approach relies greatly on the existence of highly accurate engineering models because it will be necessary to determine the mass and energy balances of each individual process unit. Third, the performance of this approach is inherently related to how well the different data silos are integrated.

# 5 Conclusion

This study highlights that the shift in energy reporting of chemical plants from an operation-based to a digital workflow intended for regulatory bodies involves a shift in paradigm with regard to the structure of such systems. This study successfully proves that a hybrid architecture in which the machine learning algorithms are controlled by the deterministic rule-based layer is effective when it comes to addressing the needs for auditability, reproducibility, and legality of the process. In terms of the results provided in this paper, it can be clearly stated that statistical analysis proved that the system proposed by the researcher had 94.8% Compliance Gain, while models relying solely on machine learning algorithms produced only 62%. In addition, the hybrid architecture proved its robustness by achieving 100% admissibility even when the prediction power of ML algorithms dropped to zero due to sensor failure. This guarantees that even when there are any missing data and/or model drifts, the generated reports still have legal soundness and physical consistency with conservation principles regarding mass and energy. The key value of the results is in the opportunity to use analytical capabilities of the artificial intelligence (e.g., anomaly detection, data reconciliation), while preserving evidential threads necessary for regulatory purposes. Future work should include the use of distributed ledger technologies (blockchain) to increase the tamper-evidencing properties of the reporting chain. Another possibility for further investigation is to consider the potential use of Large Language Models (LLMs) to translate evolving legal regulatory documents into logic rules for the rule-based layer.

# References

[1] Akbulut, O., Cavus, M., Cengiz, M., Allahham, A., Giaouris, D., & Forshaw, M. (2024). Hybrid intelligent control system for adaptive microgrid optimization: Integration of rule-based control and deep learning techniques. *Energies*, *17*(10), 2260. https://doi.org/10.3390/en17102260

[2] Alotaibi, E. M., Khallaf, A., Abdallah, A. A. N., Zoubi, T., & Alnesafi, A. (2024). Blockchain-driven carbon accountability in supply chains. *Sustainability*, *16*(24), 10872. https://doi.org/10.3390/su162410872

[3] Behzadi, A., & Sadrizadeh, S. (2023). Grid-tied solar and biomass hybridization for multi-family houses in Sweden: an optimal rule-based control framework through machine learning approach. *Renewable energy*, *218*, 119230. https://doi.org/10.1016/j.renene.2023.119230

[4] Bell, C., Ilonze, C., Duggan, A., & Zimmerle, D. (2023). Performance of continuous emission monitoring solutions under a single-blind controlled testing protocol. *Environmental science & technology*, *57*(14), 5794-5805. https://doi.org/10.1021/acs.est.2c09235

[5] Boyd, K. L. (2021). Datasheets for datasets help ML engineers notice and understand ethical issues in training data. *Proceedings of the ACM on Human-Computer Interaction*, *5*(CSCW2), 1-27. https://doi.org/10.1145/3479582

[6] Clark, L., Garcia, B., & Harris, S. (2025). Digital Twin-Driven Operational Management Framework for Real-Time Decision-Making in Smart Factories. *Journal of Innovation in Governance and Business Practices*, *1*(1), 59-87. https://doi.org/10.66096/JIGBP.V1.3

[7] Ding, N., Xi, Y., Jiang, W., Li, H., Su, J., Yang, S., & Lie, T. T. (2025). State-of-the-art carbon metering: Continuous emission monitoring systems for industrial applications. *Heliyon*, *11*(3). https://doi.org/10.1016/j.heliyon.2025.e42308

[8] Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Iii, H. D., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, *64*(12), 86-92. https://dx.doi.org/10.1145/3458723

[9] Green, W., Taylor, S., & Wu, J. (2017). Determinants of greenhouse gas assurance provider choice. *Meditari Accountancy Research*, *25*(1), 114-135. https://doi.org/10.1108/MEDAR-08-2016-0072
[10] Hino, M., Benami, E., & Brooks, N. (2018). Machine learning for environmental monitoring. *Nature Sustainability*, *1*(10), 583-588. https://doi.org/10.1038/s41893-018-0142-9
[11] Jaigirdar, F. T., Tan, B., Rudolph, C., & Bain, C. (2023). Security-aware provenance for transparency in IoT data propagation. *IEEE Access*, *11*, 55677-55691. https://doi.org/10.1109/ACCESS.2023.3280928
[12] Kierner, S., Kucharski, J., & Kierner, Z. (2023). Taxonomy of hybrid architectures involving rule-based reasoning and machine learning in clinical decision systems: A scoping review. *Journal of biomedical informatics*, *144*, 104428. https://doi.org/10.1016/j.jbi.2023.104428
[13] Lv, Z., Han, Y., Singh, A. K., Manogaran, G., & Lv, H. (2020). Trustworthiness in industrial IoT systems based on artificial intelligence. *IEEE Transactions on Industrial Informatics*, *17*(2), 1496-1504. https://doi.org/10.1109/TII.2020.2994747
[14] Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., ... & Gebru, T. (2019, January). Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency* (pp. 220-229). https://doi.org/10.1145/3287560.3287596
[15] Muñoz, L. A., Martínez, J. V. B., Pérez, F. M., & Fonseca, I. L. (2024). Anomaly detection system for data quality assurance in IoT infrastructures based on machine learning. *Internet of Things*, *25*, 101095. https://doi.org/10.1016/j.iot.2024.101095
[16] Schulte, S., Jakobs, A., & Pauliuk, S. (2024). Estimating the uncertainty of the greenhouse gas emission accounts in global multi-regional input–output analysis. *Earth System Science Data*, *16*(6), 2669-2700. https://doi.org/10.5194/essd-16-2669-2024
[17] Wagenhofer, A. (2024). Sustainability reporting: A financial reporting perspective. *Accounting in Europe*, *21*(1), 1-13. https://doi.org/10.1080/17449480.2023.2218398
[18] Wang, X., Xu, L., Zhang, Q., Zhang, D., & Zhang, X. (2022). Evaluating the data quality of continuous emissions monitoring systems in China. *Journal of environmental management*, *314*, 115081. https://doi.org/10.1016/j.jenvman.2022.115081
[19] Zhou, S. (2022). Reporting and assurance of climate-related and other sustainability information: a review of research and practice. *Australian Accounting Review*, *32*(3), 315-333. https://doi.org/10.1111/auar.12387

## Author Biography

**Maheswara Rao Gorumutchu** is associated with HYR Global Source Inc, United States. His expertise combines computer science, global business operations, and technology-driven enterprise solutions. He actively contributes to innovation-focused collaborations involving digital systems, software applications, and strategic industry development.