EAI Endorsed Transactions
on Scalable Information Systems
Research Article EAI.EU

Digital Twin and Constraint-Guided Generative Decision Support for Human-Machine Collaborative Manufacturing under Industry 5.0

Chen Qian[1,*]

[1]School of International Business, Zhejiang International Studies University, Hangzhou, 310023, China

Abstract

INTRODUCTION: In Industry 5.0 manufacturing systems, adaptive decision support is required to handle demand fluctuations, machine degradation, and operational feasibility constraints under human supervision. However, existing digital twin and data-driven approaches remain primarily focused on monitoring or forecasting, with limited integration into decision-oriented production control and feedback adaptation.
OBJECTIVES: This study aims to develop a human-in-the-loop adaptive production decision-support framework that transforms external demand signals into feasible, risk-aware, and reviewable production actions by integrating demand-production state fusion, digital twin-based constraints, structured generative decision decoding, and feedback-driven updates.
METHODS: A constrained decision-making system is constructed using Online Retail II and AI4I 2020 datasets. A unified state representation is learned by fusing demand indicators with manufacturing digital twin variables. A constraint-guided policy decoder is designed to generate discrete production actions under feasibility constraints rather than performing demand forecasting. Human feedback is encoded into subsequent decision states to enable iterative adaptation. The model is trained via a unified optimization objective balancing demand alignment, risk control, and explanation consistency.
RESULTS: Experimental results on two datasets show that the proposed framework achieves improved demand-production alignment while consistently reducing manufacturing risk exposure and human override rates compared with forecasting-based and decision-oriented baselines.
CONCLUSION: The proposed framework demonstrates that integrating digital twin constraints, constrained generative decision-making, and human feedback enables a shift from demand prediction to feasibility-aware production decision optimization, improving safety, interpretability, and human acceptability in Industry 5.0 manufacturing systems.

Keywords: industry 5.0, human-machine collaboration, digital twin, generative artificial intelligence, adaptive production decision support

Received on 29 June 2026, accepted on 21 July 2026, published on 12 August 2026

doi: 10.4108/eetsis.13777

1. Introduction

Industry 5.0 emphasizes human-centred, resilient, and adaptive manufacturing, where production systems must respond to dynamic demand, equipment risk, and human decision input while maintaining operational feasibility [1,2]. Demand fluctuations may require changes in production priority, capacity allocation, delivery scheduling, or maintenance timing [3,4]. However, such responses cannot be determined from demand information alone; feasible decisions must also consider machine condition, capacity utilization, quality deviation, failure risk, and expert judgement. Therefore, this study focuses on adaptive

*Corresponding author. Email: qianchen@zisu.edu.cn

production decision support by transforming external demand changes and manufacturing states into feasible, risk-aware, and interpretable production suggestions.

Existing studies provide useful foundations. Digital twin methods support equipment monitoring, simulation, and manufacturing-state representation [5,6]. Demand forecasting models predict order trends and variations, while Generative AI supports industrial recommendation and explanation. Human-in-the-loop manufacturing studies emphasize expert supervision. However, these approaches generally address monitoring, prediction, recommendation, or review separately rather than integrating demand variation with feasible production response.

The core gap lies in the lack of a mechanism connecting demand variation, manufacturing feasibility, constrained recommendation, and feedback-based adjustment. Forecasting methods identify demand changes but do not determine whether responses are feasible under capacity or machine-risk constraints [7]. Digital twins describe manufacturing states but are often used for monitoring rather than demand-responsive decision generation [8]. Generative AI can provide recommendations and explanations, but its outputs require operational constraints to avoid infeasible decisions. Human feedback can correct unsuitable suggestions, but without being incorporated into subsequent states, such correction remains a one-time adjustment.

To address this gap, this study proposes a human-in-the-loop adaptive production decision-support framework for Industry 5.0 manufacturing. External demand signals are integrated with digital twin states to jointly consider capacity, machine condition, quality deviation, and failure risk. A constraint-guided structured policy decoder generates production suggestions under feasibility constraints, while review feedback updates subsequent decision states. Here, generative decision refers to structured generation of production actions and explanations from fused states rather than free-form language generation by a large language model. E-commerce transaction data are used only as external demand signals rather than the central research object. The framework is mainly applicable to manufacturing scenarios with dynamic demand fluctuations, equipment-risk-sensitive decisions, and human review processes, such as order-driven and multi-variety small-batch production.

The framework includes three modules. The demand-production state fusion module connects market demand indicators with manufacturing feasibility variables. The structured policy decoder generates production adjustment, maintenance review, or human confirmation actions. The human-in-the-loop module records review outcomes, simulated override events, and revised actions in the controlled evaluation for later updates.

Different from constrained optimization with manually designed objectives, the proposed framework learns structured mappings between heterogeneous states and feasible actions. Unlike digital twin scheduling methods that mainly optimize internal processes, digital twin states are used as feasibility constraints for demand-responsive decisions. Compared with rule-based and reinforcement learning methods, the framework avoids fixed rules and reward-driven exploration through constraint-guided policy inference and feedback updates. Therefore, the proposed method serves as a decision-support framework for assisting production planners rather than replacing fully automated production-control systems.

The contributions are threefold. First, this study formulates adaptive production decision support as a demand-production alignment problem under manufacturing feasibility constraints. Second, it designs a constraint-guided structured decision mechanism linking production suggestions with capacity, risk, delay, and explanation consistency. Third, it evaluates manufacturing alignment, risk control, and simulated feasibility-review consistency using demand-production matching error, risk exposure, and simulated Human Override Rate (sHOR), with explanation consistency analyzed as an auxiliary reviewability indicator.

2. Related Works

This section reviews studies related to adaptive production decision support for Industry 5.0 manufacturing [9]. Existing research can be grouped into four streams: demand-driven production planning, digital twin-enabled manufacturing support, generative AI for industrial recommendation and explanation, and human-in-the-loop manufacturing decision-making [10]. These streams provide useful foundations, but they have not fully connected demand fluctuation, manufacturing feasibility, risk-aware recommendation, and feedback-based adjustment into an integrated production response mechanism [11].

2.1. Demand-Driven Production Planning and Forecasting

Demand-driven planning methods help enterprises anticipate market changes and adjust production or inventory decisions [12]. Traditional statistical models and recent machine learning methods are widely used to predict order volume, category-level demand, seasonal variation, and sales trends [13]. These methods transform historical transaction data into demand indicators for planning and resource preparation, while recent models better capture nonlinear demand patterns and short-term fluctuations [14].

However, demand forecasting mainly provides market-side prediction rather than manufacturing-side response. A predicted demand increase does not directly indicate whether production should be increased, delayed, or reviewed, because the decision also depends on capacity utilization, machine condition, tool wear, quality deviation, and failure risk [15]. This study differs from forecasting-oriented work by focusing on converting external demand signals into feasible production adjustment suggestions [16]. The gap is that demand-driven methods capture market variation but do not sufficiently model the feasibility mechanism behind production response.

2.2. Digital Twin-Enabled Manufacturing State Representation

Digital twin methods support smart manufacturing by representing the real-time or near-real-time state of production systems [17]. Existing studies use digital twins for equipment monitoring, production simulation, dynamic scheduling, and process optimization [18,19]. These methods improve manufacturing-state visibility by linking physical equipment with virtual representations, helping describe machine condition, process status, capacity utilization, and failure risk [17,20]. Recent studies further extend digital twins toward intelligent decision support by integrating operational data, predictive analytics, and adaptive optimization mechanisms, enabling more responsive manufacturing management [20].

Nevertheless, many digital twin studies mainly emphasize monitoring, simulation, or internal scheduling support. They do not always explain how external demand disturbances should be interpreted with digital twin states to generate concrete production adjustment suggestions. Related digital twin scheduling studies provide useful optimization mechanisms, but they usually treat production objectives and constraints as internal manufacturing variables [18,19]. This study differs by using digital twin states as feasibility constraints within a demand-responsive decision-support process. The gap is that digital twin methods provide manufacturing awareness, but the transformation from state awareness to demand-responsive and risk-aware recommendation remains underdeveloped.

2.3. Generative AI for Industrial Recommendation and Explanation

Generative AI has been applied to industrial scenarios such as scenario generation, decision recommendation, knowledge summarization, and explanation generation [21,22]. Recent research has explored large-scale generative models for manufacturing knowledge assistance and human-centered decision support, demonstrating their potential for interpreting complex industrial information and generating actionable suggestions [21,23]. Compared with conventional predictive models, such models can translate complex industrial states into readable recommendations and explanations for human-machine collaborative decision-making [24,25].

However, generative AI creates a specific risk in manufacturing applications: a recommendation may be linguistically plausible but operationally infeasible due to hallucination, insufficient grounding, or missing operational constraints [26]. Recent studies have systematically investigated hallucination and trustworthiness issues in large language models, emphasizing that generated outputs may lack factual consistency and reliable grounding even when they appear fluent and coherent [27]. If generation is not constrained by capacity, machine risk, maintenance requirements, and response delay, the model may recommend production increases simply because demand is rising, even when equipment conditions do not support them. To improve output reliability, recent work has further explored constraint decoding strategies that control black-box language-model outputs through predefined constraints rather than unconstrained generation [28].

This study differs by treating generative decision-making as a constraint-guided structured policy-decoding process rather than free-form text generation. Existing generative AI methods support recommendation and explanation, but their outputs are not always explicitly restricted by manufacturing feasibility conditions. The proposed framework addresses this gap by applying constraint projection to the candidate action space before final action selection.

2.4. Human-in-the-Loop Manufacturing Decision Support

Human-in-the-loop manufacturing research emphasizes that expert judgement remains necessary in Industry 5.0 systems [29]. Recent human-centered AI studies further highlight the importance of combining automated intelligence with human oversight to improve trust, accountability, and operational acceptance in industrial decision-making environments [30]. Human operators can assess contextual risk, correct unsuitable recommendations, and determine whether automated suggestions are acceptable in actual production conditions. These studies prevent decision support from becoming fully automated and disconnected from operational knowledge, while also improving trust, interpretability, and accountability [31].

However, human involvement is often positioned as a review or supervision step after the system generates a recommendation. Review feedback may correct the current output but is not always encoded into the next decision state. As a result, the system may repeat similar unsuitable suggestions under similar demand-risk conditions. This study differs by using review evaluation, override behaviour, and revised actions as feedback signals for later decision updates. The gap is that human-in-the-loop methods support supervision, but they do not always convert expert revision into an adaptive update mechanism.

2.5. Research Gap

Overall, existing research has advanced demand forecasting, digital twin modelling, generative AI recommendation, and human-centred manufacturing decision support. However, adaptive production response under simultaneous demand fluctuation and manufacturing-risk constraints remains insufficiently resolved. Existing methods often predict demand, represent manufacturing states, generate explanations, or involve expert review as separate functions. What is still lacking is a mechanism that can fuse external demand signals with digital twin production states, generate constrained and risk-aware production adjustment suggestions, and update later decisions using review feedback. This study addresses this gap through a human-in-the-loop adaptive production decision-support framework in which demand sensing is treated as an external disturbance input, digital twin states define manufacturing feasibility, generative AI produces constrained recommendations, and review feedback supports subsequent decision adjustment.

3. Methodology

This section presents the proposed human-in-the-loop adaptive production decision-support framework for Industry 5.0 manufacturing. The framework formulates adaptive production decision-making as a constrained policy inference problem, where external demand variations and digital twin-based manufacturing states are transformed into feasible, risk-aware, and interpretable production actions. Unlike demand forecasting studies that treat future demand as the final output, this work treats demand variation as an external disturbance input and focuses on manufacturing-side decision response under operational constraints.

3.1. Problem Definition

This study formulates the adaptive production decision-support problem as a constrained sequential decision-making task in an Industry 5.0 human-machine collaborative manufacturing environment. The objective is to learn a risk-aware policy that maps dynamic demand signals and manufacturing states into feasible production actions under operational constraints, rather than performing demand forecasting.

At each decision window t, the system observes a unified state:

$$X_t = (D_t, M_t, C_t) \tag{1}$$

where D_t denotes external demand signals, M_t represents digital twin-based manufacturing states, and C_t captures contextual constraints.

Specifically, demand variables $D_t = [s_t, g_t, r_t, c_t, v_t]$ include sales volume, demand growth rate, return ratio, category demand intensity, and volatility. Manufacturing states $M_t = [m_t, u_t, w_t, q_t, f_t]$ describe machine condition, capacity utilization, tool wear, quality deviation, and failure risk. Contextual constraints $C_t = [l_t, p_t, i_t, h_t]$ include delivery urgency, order priority, inventory level, and maintenance threshold.

The system learns a constrained policy:

$$(a_t, e_t) \sim \pi_\theta(a, e \mid X_t) \tag{2}$$

where $a_t \in A = \{A_1, \dots, A_5\}$ denotes discrete production actions, including production increase, plan maintenance, delay, maintenance triggering, and human confirmation, and e_t is the corresponding explanation generated under the same state evidence.

The feasible decision space is defined as:

$$A_t^{\text{valid}} = \{a \in A \mid u_t \le U_{\max}, f_t \le F_{\max}, w_t \le W_{\max}\} \tag{3}$$

The policy is therefore constrained as:

$$\pi_\theta \text{ is optimized over } A_t^{\text{valid}} \tag{4}$$

Historical human feedback signals further enable adaptive refinement of the policy across decision windows, supporting continuous improvement under evolving manufacturing conditions.

Samples are constructed by aligning Online Retail II demand data with AI4I 2020 manufacturing-state records within standardized decision-window indices as a controlled experimental construction. The two datasets are not collected from the same manufacturing system: Online Retail II represents external demand disturbances, while AI4I 2020 provides manufacturing-side operational states for evaluating feasibility-aware decision generation. Therefore, the alignment does not imply that the observed demand events and manufacturing states occurred at the same historical time or within the same enterprise.

Transaction-level demand events are aggregated into fixed-length decision windows to generate demand indicators, while AI4I 2020 sensor records are resampled using a corresponding number of standardized state windows through window-based aggregation. When the resulting demand and manufacturing sequences contain different numbers of windows, linear interpolation is applied to the manufacturing-state sequence over a normalized relative index so that the two sequences have the same number of controlled evaluation samples.

The resulting samples are index-aligned controlled pairs rather than naturally synchronized demand-manufacturing observations. The construction is used to test whether the proposed model can combine an external demand-disturbance representation with feasibility-related manufacturing states under controlled conditions; it does not establish real-world temporal, causal, or operational coupling between the two datasets.

The decision-window setting is used as an experimental aggregation unit rather than as evidence of an actual factory scheduling interval. The exact window length, aggregation rule, resampling procedure, and number of resulting samples are reported in Section 4.1.1 to support reproducibility.

3.2. Overall Framework

To address the disconnection between demand variation, manufacturing-state monitoring, decision recommendation, and human revision, this study proposes a closed-loop adaptive decision framework based on digital twin representation and constrained policy modeling, as illustrated in Figure 1. The system is formulated as a constrained sequential decision process governed by a parameterized policy π_θ, rather than a direct mapping function.

At each decision step t, demand variables D_t, manufacturing states M_t, and contextual constraints C_t are encoded into a unified state representation:

$$S_t = F(D_t, M_t, C_t) \tag{5}$$

where $F(\cdot)$ integrates market disturbances, machine condition, capacity utilization, and operational constraints into a shared latent embedding space.

Based on S_t, the system performs constrained policy inference over a feasibility-aware action space:

$$a_t \sim \pi_\theta(a \mid S_t, A_t^{\text{valid}}) \tag{6}$$

where A_t^{valid} denotes a dynamically constrained action space defined under capacity, risk, and maintenance constraints. The policy is implemented via structured constrained decoding over discrete production actions.

Human feedback generates a revision signal r_t, which updates the state transition dynamics:

$$S_{t+1} = g(S_t, a_t, r_t) \tag{7}$$

This forms a closed-loop constrained decision system that jointly integrates state encoding, policy-based inference, and feedback-driven adaptation, enabling continuous alignment between demand variation and manufacturing feasibility.

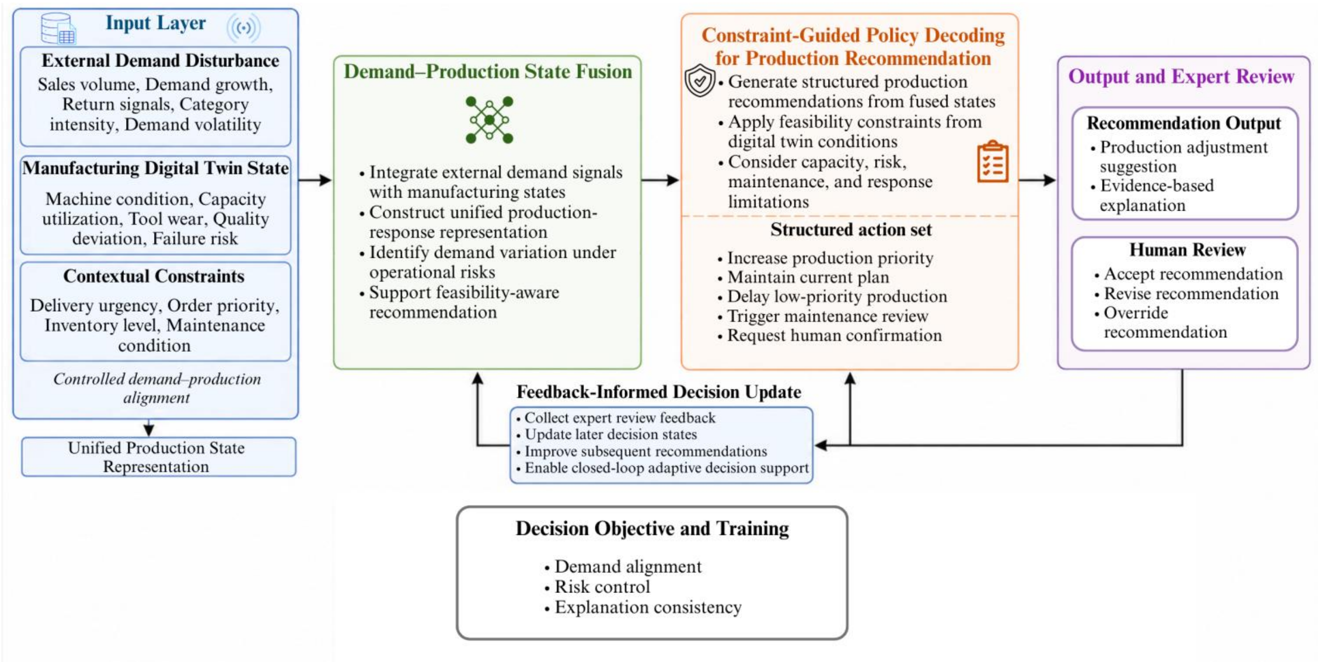

Figure 1. Overall closed-loop framework of the proposed human-in-the-loop adaptive production decision-support system

3.3. Demand-Production State Fusion

The demand-production state fusion module addresses the gap between market-side demand variation and manufacturing feasibility assessment by learning a unified decision representation that jointly encodes heterogeneous operational factors. Demand forecasting captures external growth signals but does not explicitly model feasibility under machine-state constraints. Therefore, this module constructs a shared latent embedding space for joint decision reasoning.

The input is $X_t = (D_t, M_t, C_t)$, where D_t, M_t, and C_t denote demand signals, manufacturing digital twin states, and contextual constraints, respectively. A unified representation is obtained through structured projection and nonlinear interaction modeling:

$$Z_t = \phi([W_D D_t \parallel W_M M_t \parallel W_C C_t] + b) \tag{8}$$

where W_D, W_M, W_C are learnable projection matrices, $\parallel$denotes feature concatenation, $b$is a bias term, and $\phi(\cdot)$ is a nonlinear transformation that maps heterogeneous inputs into a shared decision embedding space.

To explicitly incorporate manufacturing safety constraints, a risk-aware gating mechanism is introduced:

$$\tilde{Z}_t = Z_t \odot (1 - \rho_t) \tag{9}$$

where $\odot$ denotes element-wise modulation and ρ_t is a normalized risk coefficient defined as:

$$\rho_t = \sigma(\alpha f_t + \beta u_t + \gamma w_t) \tag{10}$$

with f_t, u_t, and w_t representing failure risk, capacity utilization, and tool wear, respectively.

Finally, a demand-risk interaction signal is defined as:

$$\kappa_t = g_t \cdot \rho_t \tag{11}$$

which captures the coupling between demand growth and system degradation, guiding conservative decision tendencies under high-conflict conditions.

3.4. Constraint-Guided Policy Decoding for Production Recommendation

The decision module is formulated as a constraint-guided structured policy decoder rather than an unconstrained text generation mechanism. It maps the fused state representation into a discrete action-explanation pair under explicit feasibility constraints defined by the digital twin environment. In this study, generative decision refers to the structured generation of production actions and associated explanations from manufacturing states rather than free-form language generation.

The input consists of the fused state representation Z_t, conflict signal κ_t, action space $\mathcal{A}$, and constraint set Ω_t. The constraint set Ω_t is constructed from digital twin-derived operational conditions, including capacity, equipment risk, tool degradation, and temporal response constraints. The decision process is modeled as a conditional policy:

$$(a_t, e_t) \sim \pi_\theta(a, e \mid Z_t, \kappa_t, \Omega_t)) \tag{12}$$

where π_θ denotes a parameterized structured policy decoder, implemented as a lightweight neural scoring function over the fused embedding space, which generates decisions over a discrete action manifold instead of free-form text generation. Therefore, the decoder serves as a structured decision inference module rather than a large language model-based text generator. The decoder consists of a lightweight two-layer neural scoring network with a 128-dimensional hidden embedding space, followed by a softmax action selection head for constrained policy inference.

To guarantee feasibility, a constraint-induced projection operator is applied to the action space:

$$\mathcal{A}_t^{\text{valid}} = \Pi_{\Omega_t}(\mathcal{A}) \tag{13}$$

where the projection operator $\Pi_{\Omega_t}(\cdot)$ eliminates infeasible actions according to capacity utilization, equipment failure risk, tool wear, and temporal response constraints derived from the digital twin system. Specifically, production-increase actions are restricted when capacity utilization, failure risk, or tool wear exceed thresholds U_{max}, F_{max}, and W_{max} , while response-infeasible actions are filtered according to L_{max}.

The detailed variable definitions, threshold values, determination criteria, and action-projection rules are summarized in Table 1. Capacity utilization, failure risk, and tool wear are represented on a normalized scale of [0,1]. In the controlled experiments, the thresholds are set to U_{max}= 0.85, F_{max} = 0.30, and W_{max} = 0.80, while the maximum allowable response delay is set to L_{max}= 7 days. These values are determined using conservative boundaries derived from the training subset distributions and validated through the validation subset and remain fixed throughout all experiments. More specifically, U_{max} defines the upper bound of allowable capacity utilization, F_{max} represents the maximum acceptable failure-risk level, W_{max} denotes the tool-wear level beyond which maintenance-related actions should be considered, and L_{max} specifies the maximum acceptable response delay.

The projection rules are manually specified by the authors for the controlled experimental setting. If $U_t \geq U_{max}$, $F_t \geq F_{max}$, or $W_t \geq W_{max}$, the production-increase action is removed from the feasible action set. When F_t or W_t exceeds its threshold, maintenance-related or human-confirmation actions remain available in the feasible action. Actions whose estimated response delay exceeds L_{max} are also removed. When multiple constraints are activated simultaneously, the final feasible action set is obtained by intersecting the action sets permitted by all activated constraints. If no feasible action remains, the decision is assigned to human confirmation. These rules are intended to reduce simulated feasibility violations in the controlled evaluation but do not constitute a formal safety guarantee for real manufacturing deployment.

Table 1. Definition of Digital Twin-Based Constraints

Constraint	Variable	Range/ unit	Threshold value	Determination basis	Exact projection rule
Capacity	U_t	[0,1]	U_{max}=0.85	Training-subset distribution analysis and validation-set verification	Remove production increase if $U_t \geq U_{max}$
Failure risk	F_t	[0,1]	F_{max}=0.30	Training-subset risk distribution and validation-set verification]	Remove expansion; retain maintenance/ human confirmation
Tool wear	W_t	[0,1]	W_{max}= 0.80	Training-subset degradation distribution and validation-set verification	Remove expansion and prioritize maintenance
Response delay	L_t	days	L_{max}= 7 days	Predefined decision-response boundary validated on validation samples	Remove actions exceeding the delay limit

The policy is therefore evaluated only over the feasible action set:

$$\pi_\theta(a \mid Z_t, \kappa_t, \Omega_t), a \in \mathcal{A}_t^{\text{valid}} \tag{14}$$

The final decision is obtained via constrained decoding:

$$a_t^* = \arg \max_{a \in \mathcal{A}_t^{\text{valid}}} \pi_\theta(a \mid Z_t, \kappa_t, \Omega_t) \tag{15}$$

The explanation e_t is generated as a conditionally coupled output of the same policy decoder:

$$e_t \sim p_\theta(e \mid a_t^*, Z_t, \kappa_t, \Omega_t) \tag{16}$$

ensuring consistency between the selected action and its associated rationale. By sharing the same state representation and constraint information, the action and explanation outputs remain aligned with the underlying manufacturing evidence.

This shared-conditioning mechanism enforces semantic alignment between decision outputs and explanatory evidence, improving coherence between production recommendations and manufacturing-state constraints under identical decision conditions.

3.5. Training Objective and Strategy

The framework is optimized under a constraint-aware policy learning objective that jointly models demand-production alignment, manufacturing risk control, explanation consistency, and response efficiency. Instead of independently optimizing heterogeneous losses, the model learns a unified constrained policy π_θ over structured decision trajectories.

The optimization objective is formulated as:

$$\mathcal{L} = \mathbb{E}_{a_t \sim \pi_\theta}\left[\mathcal{L}_{match} + \lambda_1 \mathcal{L}_{risk} + \lambda_2 \mathcal{L}_{exp} + \lambda_3 \mathcal{L}_{delay}\right] \quad (17)$$

where π_θ denotes the constrained policy defined over the feasible action space. $\mathcal{L}_{match}$ measures the discrepancy between demand variation and production response under the fused state representation. $\mathcal{L}_{risk}$ penalizes decisions associated with high failure risk f_t, capacity over-utilization u_t, and tool wear w_t, reflecting digital-twin-derived operational constraints. $\mathcal{L}_{exp}$ enforces consistency between generated explanations and decision evidence, improving traceability and interpretability. $\mathcal{L}_{delay}$ regularizes response latency τ_t to ensure real-time applicability.

The hyperparameters $\lambda_1, \lambda_2, \lambda_3$ control the trade-off among safety, interpretability, and responsiveness. Higher λ_1 induces more conservative policies under elevated risk, while λ_2 and λ_3 improve explanation quality and system responsiveness, respectively.

Training follows a structured constrained optimization procedure. First, demand and manufacturing variables are segmented into standardized index-aligned decision windows. Second, the fusion and policy modules are jointly optimized under the unified objective. Third, constrained decoding is applied over the feasible action space A_{valid} to ensure manufacturability compliance. The model does not rely on reinforcement learning; instead, it approximates a constrained policy through end-to-end structured optimization under feasibility-aware inference.

The policy decoder is trained using a supervised learning signal constructed from aligned decision tuples (X_t, a_t, e_t), where the objective includes a negative log-likelihood term $\log \pi_\theta (a_t \mid Z_t, \kappa_t, \Omega_t)$ to explicitly optimize action prediction.

3.6. Algorithm Flow and Complexity Analysis

The proposed method is formulated as a policy-driven constrained inference framework rather than stochastic search or reinforcement learning. The overall process consists of state encoding, constraint projection, and policy-based inference under a closed-loop feedback mechanism.

Given input variables D_t, M_t, C_t, the system constructs a unified state representation:

$$X_t = (D_t, M_t, C_t) \quad (18)$$

The fusion operator maps the input state into a structured decision embedding and risk signal:

$$(\tilde{Z}_t, \kappa_t) = F(X_t) \quad (19)$$

A constraint projection operator is then applied to define the feasible action space:

$$A_t^{valid} = \Pi_{\Omega_t}(A) \quad (20)$$

The constrained policy performs action selection over the feasible set:

$$a_t = \arg \max_{a \in A_t^{valid}} P_\theta(a \mid \tilde{Z}_t, \kappa_t) \quad (21)$$

The explanation is generated as a conditional policy:

$$e_t \sim \pi_\theta(e \mid a_t, \tilde{Z}_t) \quad (22)$$

Human feedback H_t is incorporated to update subsequent states, forming a closed-loop adaptive decision process.

The computational complexity mainly arises from three operations: state fusion, constraint projection, and policy decoding. Assuming $\mid A \mid$ an input feature dimension of d, a hidden representation dimension of h, and an action-space size of $\mid A \mid$, and K active feasibility constraints, the fusion module has a computational complexity related to feature interaction, approximated as $O(d^2)O(dh + h^2)$. Constraint projection requires $O(\mid A \mid)O(K \mid A \mid)$, and policy decoding requires $O(\mid A \mid d)O(h \mid A \mid)$.

Therefore, the total per-window inference complexity is approximately $O(dh + h^2 + K \mid A \mid + h \mid A \mid)$. Because the action space is fixed and relatively small in the current implementation, the empirical inference cost is mainly determined by state fusion and neural action scoring. This analysis characterizes the implemented model-side computation and does not by itself establish end-to-end real-time industrial deployability.

Algorithm 1. Constraint-Guided Adaptive Production Recommendation

```
Input: D_t, M_t, C_t, A, Ω_t
Output: a_t, e_t, X_{t+1}
Initialize policy network parameters π_θ
for each decision window t do
    Observe D_t, M_t, C_t
    Construct X_t = (D_t, M_t, C_t)
    Compute (Z_t, κ_t) = F(X_t)
    Generate constraint set Ω_t from digital twin rules
    Compute A_t^valid = Π_{Ω_t}(A)
    for each a ∈ A_t^valid do
        Compute P_θ(a | Z_t, κ_t)
    end for
    Select a_t = arg max_{a∈A_t^valid} P_θ(a | Z_t, κ_t)
    Generate e_t ~ π_θ(e | a_t, Z_t)
    Recommend a_t for review or simulated feasibility evaluation
    Collect H_t review signal H_t
    Compute L = L_match + λ_1 L_risk + λ_2 L_exp + λ_3 L_delay
    Update θ ← θ − η∇_θ L
    Update X_{t+1} = Update(X_t, a_t, H_t)
    Store (X_t, a_t, e_t, H_t)
end for
```

4. Results and Analysis

4.1. Experimental Setup

4.1.1. Dataset Setting

To evaluate adaptive production decision support rather than standalone demand forecasting, this study constructs a controlled demand-production alignment setting using two public datasets. Online Retail II represents external demand disturbances, while AI4I 2020 Predictive Maintenance represents manufacturing-side states and risk-related constraints. The two datasets are not collected from the same manufacturing system: Online Retail II provides market-side demand variations, while AI4I 2020 provides manufacturing states for feasibility-aware decision evaluation. Therefore, this setting evaluates the model's ability to process jointly constructed demand and manufacturing-state representations and constrained decision generation rather than reproducing a specific factory environment. The resulting pairs should not be interpreted as naturally synchronized or causally related demand–machine observations.

Online Retail II contains transaction records from a UK-based online retailer during 2009-2011 [32] and was obtained from the UCI Machine Learning Repository. It is used to derive demand features, including sales volume, growth rate, return-related signals, category intensity, and volatility. The dataset was originally developed for retail transaction analysis rather than manufacturing demand-production coupling; therefore, it is used in this study only as a proxy source of external demand variation. Missing key identifiers are removed where necessary, cancelled or negative orders are used to approximate return-related variables, and transactions are aggregated by category and time window. A fixed window length of one week and a stride of one week are used, producing 118 demand windows. These settings are experimental aggregation choices and are not derived from a specific factory production cycle.

AI4I 2020 is a synthetic predictive-maintenance dataset containing temperature, rotational speed, torque, tool wear, and failure labels [33] and was obtained from the UCI Machine Learning Repository. Because it is a synthetic predictive-maintenance dataset and does not contain enterprise order, production-scheduling, or planner-feedback records, it is used only to construct manufacturing-side state and risk variables in the controlled evaluation. These variables are used to construct digital twin states, including machine condition, equipment load, degradation, quality deviation, and failure risk. Numerical variables are normalized, failure labels are converted into binary risk states, and manufacturing records are aggregated into 200 manufacturing-state windows according to their original record order. Sensor variables are resampled using window averaging.

Because the datasets do not share timestamps, calendar periods, production identifiers, or enterprise provenance, direct temporal matching is not performed. Instead, both window sequences are mapped to a normalized relative index $u \in [0,1]$. When $N_m \neq N_d$, the manufacturing-state sequence is linearly interpolated to N_d index positions, after which demand window i is paired with manufacturing-state representation i. The interpolation is performed separately within the training, validation, and test partitions to avoid information leakage across data splits.

This procedure produces standardized index-aligned samples for controlled methodological evaluation. However, it does not reproduce naturally synchronized demand-manufacturing observations or establish a causal relationship between demand variation and equipment condition, which limits direct generalization to real industrial environments. The constructed sample pairs are divided in an order-preserving manner into 70% training, 15% validation, and 15% testing subsets. All normalization, resampling, and interpolation parameters are determined using the training subset and then applied to the validation and testing subsets to avoid information leakage.

4.1.2. Baselines

As summarized in Table 2, four forecasting baselines are selected. ARIMA represents statistical time-series forecasting [34], LSTM represents deep sequential forecasting [35], Sparse Transformer with Local and Seasonal Adaptation represents transformer-based forecasting [36], and the Hybrid Demand Forecasting Model represents ensemble forecasting [37].

To provide both forecasting-oriented and decision-oriented references, a rule-based production scheduling policy [38] and a reinforcement learning (RL) policy [39] are additionally included. The rule-based method maps demand and capacity into fixed actions, while the RL policy learns sequential production decisions under reward constraints.

Table 2. Baseline Methods Used in the Experiments

Type	Method	Reason
Classical	ARIMA	Statistical forecasting
Classical	LSTM	Sequential forecasting
Strong	Sparse Transformer	Multivariate forecasting
Strong	Hybrid Forecasting Model	Ensemble forecasting
Decision	Rule-based Policy	Heuristic scheduling
Decision	RL Policy	Learned scheduling policy

For comparison, forecasting baselines are converted into production actions using a common mapping procedure. Let $\widehat{g_t}$denote the demand-change value predicted by a forecasting baseline for decision window t. A predicted demand increase above the upper threshold $\delta_+ =$ 0.10 is mapped to the "increase production" action. A predicted demand change within the stable interval $[\delta_-, \delta_+]$,where $\delta_- =$ -0.10, is mapped to "maintain the current production plan." A predicted demand decrease below δ_- is mapped to is mapped to delay low-priority production.

These forecasting-to-action mapping rules are manually defined by the authors for the controlled decision-level comparison and are not derived from an external production-planning standard. The thresholds are determined using the 10th and 90th percentiles of demand-change distributions in the training subset and remain fixed across all forecasting baselines and test samples.

After the preliminary action is generated, the same capacity, failure-risk, tool-wear, and response-delay projection rules described in Section 3.4 are applied. If the mapped action violates one or more manufacturing constraints, it is replaced by human confirmation. Therefore, all forecasting baselines use the same action categories, threshold values, and feasibility-filtering procedure. The forecasting conditions, preliminary actions, and constraint-handling rules are summarized in Table 3.

Table 3. Unified Forecasting-to-Action Mapping Rules for Forecasting Baselines

Forecast condition	Preliminary action	Feasibility treatment
$\widehat{g_t} > \delta_+$	Increase production	Project onto the feasible action set defined in Section 3.4
$\delta_- < \widehat{g_t} < \delta_+$	Maintain the current production plan	Retain unless another manufacturing constraint is triggered
$\widehat{g_t} < \delta_-$	Delay low-priority production	Apply the same constraint-projection rules
Mapped action violates one or more constraints	—	Replace with human confirmation action

Because the forecasting baselines estimate demand variation rather than directly learning feasibility-aware actions, they are treated as forecasting-oriented reference methods rather than architecturally equivalent decision models.

4.1.3. Implementation Details

All experiments are implemented using Python and PyTorch on a workstation equipped with an NVIDIA GeForce RTX 4090 GPU, CUDA 11.7, and 64 GB RAM. The batch size is 32, the optimizer is AdamW, the initial learning rate is 1×10^{-4}, and the weight decay is 1×10^{-4}. The maximum training epoch is 1000, with early stopping based on validation DPME. Rolling windows are used to compute demand growth and volatility. Each experiment is repeated three times with different random seeds, and results are reported as mean ± standard deviation.

4.1.4. Evaluation Metrics

Three metrics are used to evaluate demand alignment, manufacturing safety, and simulated feasibility-review consistency. Demand-production matching error (DPME) measures the gap between demand variation and production response, defined as

$$\mathrm{DPME} = \frac{1}{T}\sum_{t=1}^{T} \| D_t - \widehat{P}_t \|_2, \tag{23}$$

where D_t and $\widehat{P}_t$ denote demand and generated production signals. Risk Exposure (RE) evaluates whether production-increase actions are generated under high-risk machine states, defined as

$$\mathrm{RE} = \frac{1}{T}\sum_{t=1}^{T} \mathbb{I}(a_t = a^{inc}) \cdot \mathbb{I}(f_t > \tau_f), \tag{24}$$

where f_t is failure risk and τ_f is the risk threshold.

Simulated Human Override Rate (sHOR) measures the proportion of recommendations that are revised or rejected by an author-defined feasibility-rule-based review procedure, defined as

$$\mathrm{sHOR} = \frac{1}{T}\sum_{t=1}^{T} \mathbb{I}(a_t \neq a_t^{review}). \tag{25}$$

Since public datasets provide no real expert intervention records, sHOR is calculated through an author-defined feasibility-rule-based simulation rather than real planner or expert-participant experiments. Each recommendation is evaluated using predefined feasibility criteria, including capacity, machine risk, tool degradation, and maintenance requirements. Recommendations violating these criteria are modified or rejected and counted as override events. The review rules are manually defined by the authors for the controlled experimental setting and are not derived from observed planner behavior, real expert annotations, or an external industrial standard. Therefore, sHOR should be interpreted as a simulated rule-consistency indicator rather than a direct measure of actual human acceptance. Lower DPME, RE, and sHOR values indicate better demand-production alignment, lower risk exposure, and greater consistency with the predefined feasibility-review rules, respectively.

4.2. Main Performance Comparison

Table 4 compares the proposed framework with six baselines on the two datasets, including four forecasting-based methods and two decision-oriented baselines. The results are reported as mean ± standard deviation over three runs. The comparison focuses on three metrics: DPME for demand-production alignment, Risk Exposure for manufacturing safety, and sHOR for consistency with predefined feasibility-review rules. To provide a clearer visual comparison, Figure 2 further presents the performance differences among all methods across the three metrics.

Table 4. Main Results on Two Datasets (mean ± SD, n = 3)

Dataset	Method	DPME ↓	Risk Exposure ↓	sHOR ↓
Online Retail II	ARIMA	0.173 ± 0.012	0.168 ± 0.013	0.259 ± 0.020
Online Retail II	LSTM	0.151 ± 0.010	0.152 ± 0.012	0.227 ± 0.016
Online Retail II	Sparse Transformer	0.139 ± 0.009	0.144 ± 0.011	0.212 ± 0.017
Online Retail II	Hybrid Ensemble	0.134 ± 0.008	0.136 ± 0.010	0.205 ± 0.016
Online Retail II	Rule-based Policy	0.163 ± 0.011	0.160 ± 0.013	0.242 ± 0.019
Online Retail II	RL Policy	0.146 ± 0.010	0.147 ± 0.013	0.219 ± 0.017
Online Retail II	Proposed	0.128 ± 0.006	0.131 ± 0.010	0.189 ± 0.014
AI4I 2020	ARIMA	0.166 ± 0.013	0.182 ± 0.015	0.272 ± 0.020
AI4I 2020	LSTM	0.148 ± 0.010	0.155 ± 0.013	0.230 ± 0.018
AI4I 2020	Sparse Transformer	0.136 ± 0.009	0.140 ± 0.012	0.209 ± 0.016
AI4I 2020	Hybrid Ensemble	0.129 ± 0.008	0.128 ± 0.012	0.204 ± 0.015
AI4I 2020	Rule-based Policy	0.158 ± 0.012	0.163 ± 0.014	0.237 ± 0.020
AI4I 2020	RL Policy	0.142 ± 0.009	0.144 ± 0.012	0.215 ± 0.017
AI4I 2020	Proposed	0.132 ± 0.007	0.116 ± 0.011	0.193 ± 0.015

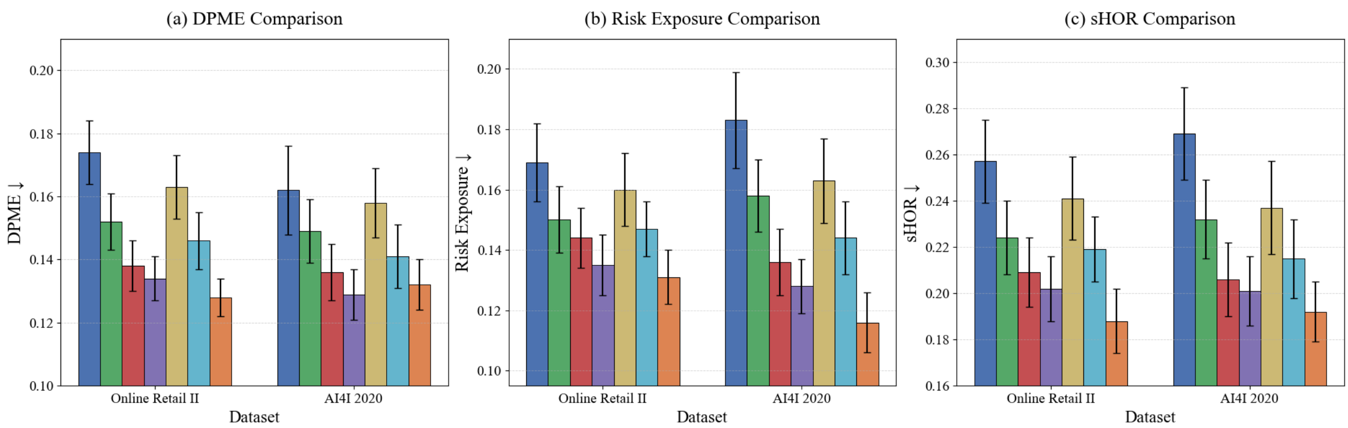

Figure 2. Main Performance Comparison on Online Retail II and AI4I 2020

(a) DPME comparison; (b) Risk Exposure comparison; (c)sHOR comparison. Lower values indicate better performance

As shown in Table 3 and Figure 2, the proposed framework achieves strong performance across both datasets, particularly in Risk Exposure and Override Rate. On Online Retail II, the proposed method achieves the lowest DPME (0.128 ± 0.006), indicating improved conversion from demand variation to feasibility-aware production decisions. Compared with Hybrid Ensemble, Risk Exposure decreases from 0.136 to 0.131, and Override Rate decreases from 0.205 to 0.189. On AI4I 2020, although Hybrid Ensemble achieves slightly lower DPME, the proposed framework achieves the best performance in both Risk Exposure and Override Rate. This indicates that the method is not purely a forecasting model, but a constraint-aware decision-making framework prioritizing lower risk exposure and greater consistency with predefined feasibility-review rules. The decision-oriented baselines (rule-based and RL policies) show competitive performance under certain conditions; however, they do not simultaneously optimize demand alignment, manufacturing risk control, and human-oriented review requirements. Although Hybrid Ensemble achieves slightly better DPME on AI4I 2020, this metric only reflects demand-production matching accuracy. The proposed framework provides additional advantages in Risk Exposure and Override Rate by explicitly considering manufacturing constraints and the author-defined feasibility-review rules.

Overall, results validate that demand-production fusion improves alignment, digital twin constraints reduce unsafe actions, and constrained policy decoding with feedback improves decision reliability. The slight DPME trade-off reflects a shift from prediction accuracy to safe decision optimization.

4.3. Statistical Significance Analysis

To examine whether the observed differences are statistically reliable, paired t-tests were conducted between the proposed framework and the strongest baseline under each dataset setting. Table 5 reports the p-values and 95% confidence intervals (CI) for the differences between the proposed framework and the strongest baseline. Since all experiments were repeated three times, the test was performed on the three repeated results of each method. The significance level was set to α=0.05, and 95% confidence intervals were reported for the performance difference. For all three metrics, lower values indicate better performance. The strongest baseline is selected as Hybrid Ensemble because it achieves the best overall performance among forecasting-based methods, ensuring a conservative statistical comparison setting.

Table 5. Statistical Significance Compared with the Strongest Baseline

Dataset	Metric	Strongest Baseline	Difference	95% CI	p-value
Online Retail II	DPME ↓	Hybrid Ensemble	-0.006	[-0.011, -0.001]	0.032
Online Retail II	Risk Exposure ↓	Hybrid Ensemble	-0.006	[-0.012, -0.001]	0.041
Online Retail II	**sHOR ↓**	Hybrid Ensemble	-0.016	[-0.025, -0.006]	0.018
AI4I 2020	DPME ↓	Hybrid Ensemble	+0.003	[-0.004, 0.010]	0.214
AI4I 2020	Risk Exposure ↓	Hybrid Ensemble	-0.011	[-0.019, -0.003]	0.041
AI4I 2020	**sHOR ↓**	Hybrid Ensemble	-0.011	[-0.020, -0.002]	0.036

As shown in Table 3, the proposed framework achieves statistically significant improvements on all three metrics for Online Retail II. On AI4I 2020, the improvement in DPME is not statistically significant because the Hybrid Ensemble baseline obtains a slightly lower matching error. However, the proposed framework shows significant reductions in Risk Exposure and Override Rate. This indicates that its advantage is concentrated on manufacturing-risk control and review-based acceptability rather than pure demand-production matching under all conditions. The confidence intervals for Risk Exposure and Override Rate are consistently below zero, suggesting that the reductions are stable across repeated runs. This result supports the core claim that digital-twin constraints and constrained policy-based recommendation can reduce risky production suggestions and improve human-machine decision acceptability.

4.4. Ablation Study

To verify the contribution of each core mechanism, ablation experiments were conducted on the two datasets under the same preprocessing, data split, and training settings. Table 6 reports the results of removing one component at a time. All results are reported as mean ± standard deviation over three runs. The evaluated variants include removing demand-production fusion, risk constraint, constrained policy decoder, and human feedback.

Table 6. Ablation Study Results on Two Datasets (mean ± SD, n = 3)

Dataset	Variant	DPME ↓	Risk Exposure ↓	sHOR ↓
Online Retail II	Full Proposed	0.128 ± 0.006	0.131 ± 0.009	0.188 ± 0.014
Online Retail II	w/o Demand-Production Fusion	0.147 ± 0.008	0.139 ± 0.011	0.211 ± 0.016
Online Retail II	w/o Risk Constraint	0.135 ± 0.007	0.153 ± 0.012	0.207 ± 0.015
Online Retail II	w/o Constrained Policy Decoder	0.138 ± 0.008	0.148 ± 0.011	0.224 ± 0.017
Online Retail II	w/o Human Feedback	0.136 ± 0.007	0.142 ± 0.010	0.246 ± 0.019
AI4I 2020	Full Proposed	0.132 ± 0.007	0.116 ± 0.010	0.192 ± 0.014
AI4I 2020	w/o Demand-Production Fusion	0.149 ± 0.009	0.132 ± 0.011	0.218 ± 0.016
AI4I 2020	w/o Risk Constraint	0.138 ± 0.008	0.151 ± 0.013	0.231 ± 0.018
AI4I 2020	w/o Constrained Policy Decoder	0.140 ± 0.008	0.143 ± 0.012	0.238 ± 0.019
AI4I 2020	w/o Human Feedback	0.137 ± 0.008	0.136 ± 0.011	0.268 ± 0.021

Removing demand-production fusion leads to the largest degradation in DPME, indicating that demand signals cannot be effectively mapped into feasible production actions without manufacturing-state representation. Removing risk constraint results in the highest Risk Exposure, demonstrating the necessity of digital twin-based safety constraints for preventing unsafe production-increase decisions. Removing structured policy decoding increases Override Rate, suggesting that unconstrained decision generation reduces alignment with expert-rule-based review. Removing human feedback further increases Override Rate, confirming the importance of feedback-driven adaptation in reducing repeated suboptimal decisions.

To evaluate the computational overhead of the constraint-guided policy decoder, inference latency was measured for all methods under the hardware and software environment reported in Section 4.1.3. As shown in Table 7, the proposed framework introduces additional computation from state fusion, constraint projection, and policy decoding compared with lightweight forecasting baselines. However, the average model-side inference latency is 38.6 ± 1.0 ms per decision window, remaining below 40 ms under the tested environment.

Latency was measured using a batch size of 1, after 20 warm-up runs, and averaged over 100 repeated forward passes. GPU synchronization was performed before and after each measurement to ensure accurate latency recording. The reported values include model forward inference and action-space projection but exclude data acquisition, file or database access, network transmission, preprocessing outside the model pipeline, platform communication, and operator interaction. Therefore, the results characterize model-side computational overhead rather than complete end-to-end industrial deployment latency.

Table 7. Model-Side Inference Latency and Additional Operations

Method	Model-side inference latency (ms/decision window) ↓	Additional Operations
ARIMA	12.4 ± 0.3	Time-series forecasting
LSTM	18.7 ± 0.5	Temporal feature encoding
Sparse Transformer	26.5 ± 0.7	Attention-based forecasting
Hybrid Ensemble	31.2 ± 0.8	Multi-model prediction fusion
Proposed Framework	38.6 ± 1.0	State fusion + constraint projection + policy decoding

4.5. Decision-State Space Visualization

To examine whether the fusion mechanism improves the separability of different decision patterns in the learned representation space, fused decision states were projected into two dimensions using t-SNE. As shown in Figure 3, each point represents one decision-window sample, grouped as feasible response, high-risk production pressure, or human-confirmation case.

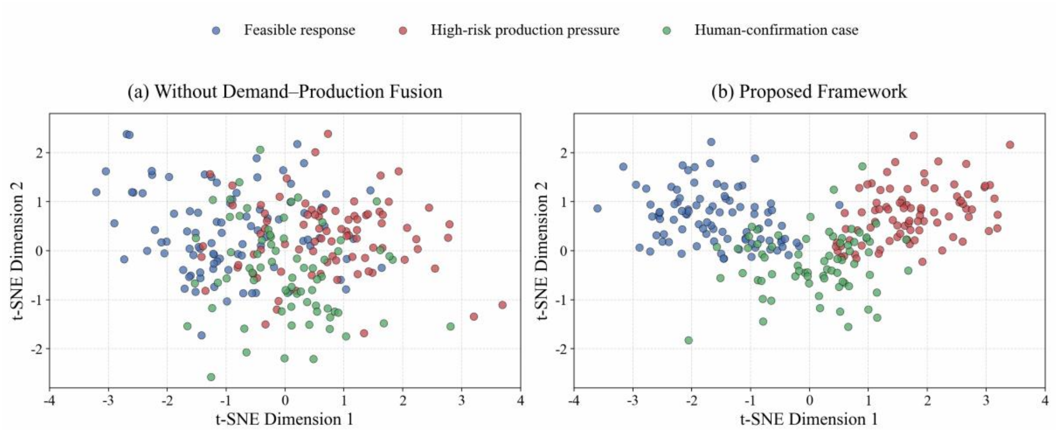

Figure 3. Decision-State Space Visualization of Fused Production Decision States. Each point represents a decision-window sample projected into a two-dimensional representation space, where different groups indicate different production decision patterns. The figure compares the separability of learned decision states before and after demand-production fusion. (a) Without demand-production fusion; (b) proposed framework

Figure 3 shows that, without demand-production fusion, the learned decision representations exhibit greater overlap between demand-growth samples and high-risk machine-state samples. In contrast, the proposed framework shows clearer separation between feasible responses and high-risk production pressure. This is consistent with the constraint-aware fusion mechanism defined in Section 3.3, where the manufacturing risk coefficient and demand-risk conflict score adjust the fused decision representation.

From a manufacturing perspective, this separation means that the framework better distinguishes "demand growth with feasible capacity" from "demand growth under unsafe machine conditions." Human-confirmation samples are mainly located between feasible and high-risk regions, suggesting that ambiguous demand-risk conflicts are more likely assigned to the author-defined feasibility-review condition. Together with the ablation results in Section 4.4, Figure 3 provides mechanism-level evidence that demand-production fusion improves representation separability and supports more risk-aware production decision support.

4.6. Explanation Consistency and Simulated Feasibility-Review Analysis

To evaluate whether the generated recommendations are understandable and traceable and consistent with predefined feasibility-review rules for human-machine collaborative manufacturing, this section analyzes explanation consistency, evidence coverage, and sHOR. Unlike attention-based visualization, this analysis focuses on whether the generated explanation is supported by demand, capacity, and risk evidence used in the decision process. Figure 4 compares the proposed framework with the strongest baseline and two ablated variants.

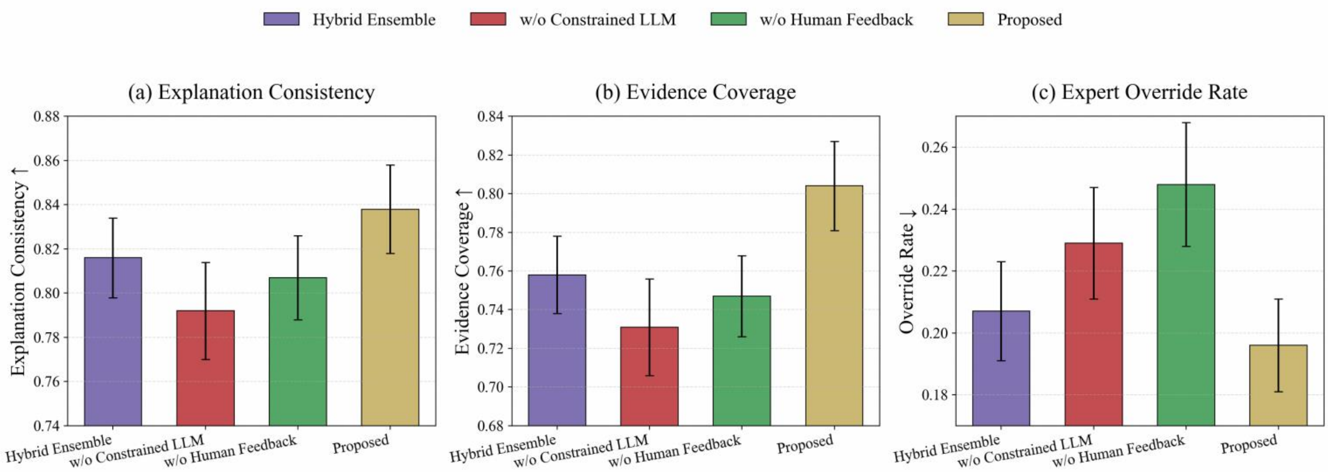

Figure 4. Explanation Consistency, Evidence Coverage, and Simulated Feasibility-Review Comparison. The figure compares the proposed framework with the strongest baseline and two ablated variants. (a) Explanation consistency measures agreement between the selected action and its generated rationale; (b) evidence coverage measures whether demand, capacity, risk, and tool-wear evidence are reflected in the explanation; and (c) sHOR reports the proportion of recommendations revised or rejected by the author-defined feasibility-review rules.

As shown in Figure 4, the proposed framework achieves higher explanation consistency and evidence coverage than the strongest baseline and the ablated variants. This indicates that the generated explanations are more closely aligned with the demand, capacity, and risk evidence used for action selection. In manufacturing terms, the system does not only state what action should be taken, but also provides traceable reasons, such as demand growth, high tool wear, or elevated failure risk.

The lower override rate further suggests that the expert-rule-based review mechanism is less likely to revise or reject the generated recommendations. Compared with the variant without structured policy decoding, the proposed framework reduces override rate because infeasible actions are filtered before constrained decision generation. Compared with the variant without human feedback, the proposed framework produces more acceptable suggestions in later decision windows. These results support the role of constraint-guided structured decision recommendation and feedback-informed update in improving the reviewability and acceptance of production decision support.

5. Discussion

This study proposes a human-in-the-loop adaptive production decision-support framework for Industry 5.0 manufacturing. Compared with traditional demand forecasting approaches such as ARIMA and LSTM-based models [27,28], which focus on predicting future demand signals, the proposed method shifts toward constrained decision generation under manufacturing feasibility conditions. Unlike forecasting and scheduling methods that improve prediction accuracy or optimize policies under uncertainty [29,32], our framework integrates demand variation with digital twin states to support risk-aware production decisions under explicit manufacturing feasibility constraints.

Digital twin-based manufacturing studies [17] mainly focus on system monitoring and simulation but rarely establish closed-loop connections among demand signals, decision generation, and human feedback. In contrast, this study introduces feedback-driven state updates to enable continuous adaptation in dynamic manufacturing environments and reduce simulated feasibility violations and repeated unsuitable recommendations.

Generative AI-based industrial methods [20,24] provide recommendation and explanation capabilities but may lack explicit operational feasibility constraints. In this study, generative decision refers to the structured generation of discrete production actions and evidence-linked explanations rather than free-form text generation. The proposed framework addresses this issue by embedding constraint projection into policy decoding, restricting action selection to a predefined feasible action set while maintaining interpretability and traceability for human decision-makers.

Overall, integrating digital twin modeling, constraint-guided structured decision generation, and human feedback provides an operationally consistent decision-support paradigm for Industry 5.0 manufacturing systems. The observed performance is mainly associated with the explicit combination of demand signals and machine-state constraints, which reduces ambiguity in action selection under the controlled experimental setting. The framework is conceptually intended for scenarios with dynamic order fluctuations, equipment-risk-sensitive production requirements, and human review processes, such as order-driven and multi-variety small-batch manufacturing. It is intended to provide reviewable recommendations for production planners rather than replace fully automated production-control systems.

For practical deployment, the decision-window length should be configured jointly with the production takt time, equipment-monitoring frequency, order-update frequency, and planner response cycle. Shorter windows may improve responsiveness but can amplify transient fluctuations and missing-data effects, whereas longer windows provide more stable state estimates but may delay production adjustment. Therefore, the window configuration should be calibrated using factory-specific operational data rather than directly adopting the controlled experimental setting used in this study.

Real-world implementation requires timestamp synchronization, buffering, and data-quality monitoring for heterogeneous sampling rates, delays, and missing values. Demand, machine-state, and production-execution records should share a common timestamp reference, while asynchronously arriving observations can be buffered within each decision window. Short data gaps may be handled through interpolation or the most recent valid observation, whereas windows with extensive missingness or abnormal sensor noise should be flagged for manual review.

The lightweight policy decoder and fixed action space limit model-side inference overhead. Under the reported hardware setting, the average model inference time is 38.6 ms per decision window. However, this measurement excludes sensor acquisition, network transmission, database access, preprocessing, platform communication, and operator interaction. Therefore, the reported latency indicates potential suitability for human-in-the-loop decision support under the tested environment rather than demonstrating complete real-time industrial deployment.

In a practical deployment, planner feedback could be recorded through approval, modification, and rejection actions and stored as review labels for subsequent offline model updating. Because the present study does not involve real planner participation, this mechanism should be regarded as a proposed implementation pathway. The framework could be connected with ERP systems for demand information, MES/APS platforms for production and scheduling data, and digital twin or industrial IoT platforms for equipment-state monitoring. A planner-facing interface could present recommended actions, supporting evidence, and triggered constraints. The framework should operate as an advisory layer rather than directly issuing autonomous production-control commands.

Because the current evaluation relies on public datasets and offline experiments, future pilot deployment should assess end-to-end latency, data drift, interface reliability, cybersecurity, planner workload, and failure-handling procedures using naturally synchronized demand, equipment-state, and production records from a real manufacturing environment.

6. Conclusion

This study addressed adaptive production decision support under dynamic demand and manufacturing-risk constraints. A human-in-the-loop framework was developed to integrate demand signals, digital twin states, constrained policy decoding, and review feedback into a unified decision process.

Through demand-production fusion, the framework aligns market variation with manufacturing feasibility. Through constrained policy decoding, it restricts production-action selection to a predefined feasible action set under capacity, risk, and maintenance constraints. Through feedback updates, review signals are incorporated into subsequent decision states for continuous adaptation.

Experiments on two datasets show that the proposed framework achieves competitive demand-production alignment while improving risk control and consistency with predefined feasibility-review rules compared with forecasting, rule-based, and reinforcement learning baselines. It consistently reduces risk exposure and simulated Human Override Rate (sHOR).

Limitations include reliance on aligned datasets rather than real industrial systems and predefined constraint rules. Therefore, the current evaluation should be interpreted as a controlled validation of jointly constructed demand and manufacturing-state processing under constraint-guided decision generation rather than a direct replication of a specific factory environment. The proposed framework is mainly applicable to manufacturing scenarios with available demand signals, digital twin-based operational states, and human-in-the-loop decision processes, such as order-driven production, risk-sensitive manufacturing, and multi-variety production systems.

Future work will further investigate real enterprise deployment, including adaptive decision-window configuration, industrial data synchronization, integration with ERP/MES/APS platforms, and more comprehensive human feedback collection mechanisms. For practical pilot deployment, edge devices could support sensor-data acquisition, preliminary preprocessing, and lightweight inference, while cloud resources could provide model updating, data storage, and large-scale decision analysis. Such deployment would also require factory-specific hardware configuration, communication interfaces, timestamp synchronization, data-quality monitoring, and operator-review support.

This edge-cloud architecture is a proposed deployment pathway rather than a configuration validated in the current experiments. Future pilot studies should therefore evaluate end-to-end latency, edge-device computation and memory requirements, communication reliability, interface compatibility, and planner interaction using naturally synchronized demand, equipment-state, and production records from a real manufacturing environment.

References

[1] Ghobakhloo M, Mahdiraji HA, Iranmanesh M, Jafari-Sadeghi V. From Industry 4.0 digital manufacturing to Industry 5.0 digital society: A roadmap toward human-centric, sustainable, and resilient production. *Information Systems Frontiers*. 2024:1-33. https://doi.org/10.1007/s10796-024-10476-z

[2] Martini B, Bellisario D, Coletti P. Human-centered and sustainable artificial intelligence in industry 5.0: Challenges and perspectives. *Sustainability*. 2024;16(13):5448. https://doi.org/10.3390/su16135448

[3] Cai Y, He Y, Shi R, Feng T, Li J. Resilience-oriented approach of dynamic production and maintenance scheduling optimisation considering operational uncertainty. *International Journal of Production Research*. 2024;62(21):7812-7835. https://doi.org/10.1080/00207543.2024.2329324

[4] Figueroa AJ, Poler R, Andres B. Adaptive production rescheduling system for managing unforeseen disruptions. *Mathematics*. 2024;12(22):3478. https://doi.org/10.3390/math12223478

[5] Zhang C, Juraschek M, Herrmann C. Deep reinforcement learning-based dynamic scheduling for resilient and sustainable manufacturing: A systematic review. *Journal of Manufacturing Systems*. 2024;77:962-989. https://doi.org/10.1016/j.jmsy.2024.10.026

[6] Cai Y, He Y, Shi R, Liao R, Cao H, Guo H, Lu H. Resilience-oriented adaptive predictive maintenance optimization for continuous process manufacturing systems considering mission profile variation. *Computers & Industrial Engineering*. 2024;197:110532. https://doi.org/10.1016/j.cie.2024.110532

[7] Babazadeh R, Taraghi Nazloo H, Kamran MA. A hybrid ANN-MILP model for agile recovery production planning for PPE products under sharp demands. *International Journal of Production Research*. 2025;63(2):758-778. https://doi.org/10.1080/00207543.2024.2313100

[8] Lee H. Fault-Resilient Manufacturing Scheduling with Deep Learning and Constraint Solvers. *Applied Sciences*. 2025;15(4):1771. https://doi.org/10.3390/app15041771

[9] Douaioui K, Oucheikh R, Benmoussa O, Mabrouki C. Machine learning and deep learning models for demand forecasting in supply chain management: A critical review. *Applied System Innovation*. 2024;7(5):93. https://doi.org/10.3390/asi7050093

[10] Ahangar MN, Farhat ZA, Sivanathan A. AI trustworthiness in manufacturing: challenges, toolkits, and the path to industry 5.0. *Sensors*. 2025;25(14):4357. https://doi.org/10.3390/s25144357

[11] Bussieweke F, Mula J, Campuzano-Bolarin F. Optimisation of recovery policies in the era of supply chain disruptions: a system dynamics and reinforcement learning approach. *International Journal of Production Research*. 2025;63(5):1649-1673. https://doi.org/10.1080/00207543.2024.2383293

[12] Elyasi M, Altan B, Ekici A, Özener ÖÖ, Yanıkoğlu İ, Dolgui A. Production planning with flexible manufacturing systems under demand uncertainty. *International Journal of Production Research*. 2024;62(1-2):157-170. https://doi.org/10.1080/00207543.2023.2288722

[13] Shao Z, Li W, Tan Y, Otto K. A systematic energy-aware scheduling framework for manufacturing factories integrated with renewables. *International Journal of Production Research*. 2024;62(21):7644-7659. https://doi.org/10.1080/00207543.2024.2323067

[14] Chowdhury AR, Paul R, Rozony FZ. A systematic review of demand forecasting models for retail e-commerce enhancing accuracy in inventory and delivery planning. *International Journal of Scientific Interdisciplinary Research*. 2025;6(1):01-27. https://doi.org/10.63125/mbbfw637

[15] Holguin Jimenez S, Trabelsi W, Sauvey C. Multi-objective production rescheduling: A systematic literature review. *Mathematics*. 2024;12(20):3176. https://doi.org/10.3390/math12203176

[16] Xu H, Zheng J, Huang L, Tao J, Zhang C. Solving dynamic multi-objective flexible job shop scheduling problems using a dual-level integrated deep Q-network approach. *Processes*. 2025;13(2):386. https://doi.org/10.3390/pr13020386

[17] Hu S, Li C, Li B, Yang M, Wang X, Gao T, et al. Digital twins enabling intelligent manufacturing: from methodology to application. *Intelligent and Sustainable Manufacturing*. 2024;1(1):10007. https://doi.org/10.35534/ism.2024.10007

[18] Ouahabi N, Chebak A, Kamach O, Laayati O, Zegrari M. Leveraging digital twin into dynamic production scheduling: A review. *Robotics and Computer-Integrated Manufacturing*. 2024;89:102778. https://doi.org/10.1016/j.rcim.2024.102778

[19] Grznár P, Papánek L, Marčan M, Krajčovič M, Antoniuk I, Mozol Š, Mozolová L. Enhancing production efficiency through digital twin simulation scheduling. *Applied Sciences*. 2025;15(7):3637. https://doi.org/10.3390/app15073637

[20] Webb L, Tokhi OM, Alkan B. State of the art and future directions of digital twin-enabled smart assembly automation in discrete manufacturing industries. *International Journal of Computer Integrated Manufacturing*. 2025;38(8):1126-1160. https://doi.org/10.1080/0951192x.2024.2387775

[21] Ghobakhloo M, Fathi M, Iranmanesh M, Vilkas M, Grybauskas A, Amran A. Generative artificial intelligence in manufacturing: opportunities for actualizing Industry 5.0 sustainability goals. *Journal of Manufacturing Technology

Management. 2024;35(9):94-121. https://doi.org/10.1108/jmtm-12-2023-0530

[22] Filz MA, Thiede S. Generative AI in manufacturing systems: Reference framework and use cases. *Procedia CIRP*. 2024;130:238-243. https://doi.org/10.1016/j.procir.2024.10.082

[23] Li Y, Zhao H, Jiang H, Pan Y, Liu Z, Wu Z, et al. Large language models for manufacturing. *Journal of Manufacturing Systems*. 2026;86:516-545. https://doi.org/10.1016/j.jmsy.2026.02.014

[24] Jackson I, Ivanov D, Dolgui A, Namdar J. Generative artificial intelligence in supply chain and operations management: a capability-based framework for analysis and implementation. *International Journal of Production Research*. 2024;62(17):6120-6145. https://doi.org/10.1080/00207543.2024.2309309

[25] Konstantinou C, Antonarakos D, Angelakis P, Gkournelos C, Michalos G, Makris S. Leveraging Generative AI Prompt Programming for Human-Robot Collaborative Assembly. *Procedia CIRP*. 2024;128:621-626. https://doi.org/10.1016/j.procir.2024.03.040

[26] Handler A, Larsen KR, Hackathorn R. Large language models present new questions for decision support. *International Journal of Information Management*. 2024;79:102811. https://doi.org/10.1016/j.ijinfomgt.2024.102811

[27] Lin Z, Guan S, Zhang W, Zhang H, Li Y, Zhang H. Towards trustworthy LLMs: a review on debiasing and dehallucinating in large language models. *Artificial Intelligence Review*. 2024;57(9):1.

[28] Geng S, Döner B, Wendler C, Josifoski M, West R. Sketch-guided constrained decoding for boosting blackbox large language models without logit access. In: *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. 2024 Aug. p. 234-245.

[29] Lee JY. Introduction of human-in-the-loop in smart manufacturing (H-SM). *International Journal of Precision Engineering and Manufacturing-Smart Technology*. 2024;2(2):209-214. https://doi.org/10.57062/ijpem-st.2024.00115

[30] Passalacqua M, Pellerin R, Magnani F, Doyon-Poulin P, Del-Aguila L, Boasen J, Léger PM. Human-centred AI in industry 5.0: a systematic review. *International Journal of Production Research*. 2025;63(7):2638-2669. https://doi.org/10.1080/00207543.2024.2406021

[31] Kim DB, Bajestani MS, Lee JY, Shin SJ, Kim GY, Sajadieh SMM, Noh S. Human-in-the-loop in smart manufacturing (H-SM): A review and perspective. *Journal of Manufacturing Systems*. 2025;82:178-199. https://doi.org/10.1016/j.jmsy.2025.05.020

[32] Chen D. Online Retail II [Dataset]. UCI Machine Learning Repository; 2012. https://doi.org/10.24432/C5CG6D

[33] Matzka S. Explainable artificial intelligence for predictive maintenance applications. In: *2020 Third International Conference on Artificial Intelligence for Industries (AI4I)*. IEEE; 2020 Sep. p. 69-74. https://doi.org/10.1109/ai4i49448.2020.00023

[34] Makridakis S, Hibon M. ARMA models and the Box-Jenkins methodology. *Journal of Forecasting*. 1997;16(3):147-163. https://doi.org/10.1002/(SICI)1099-131X(199705)16:3<147::AID-FOR653>3.0.CO;2-X

[35] Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Computation*. 1997;9(8):1735-1780. https://doi.org/10.1162/neco.1997.9.8.1735

[36] Zhang Y, Wu R, Dascalu SM, Harris Jr FC. Sparse transformer with local and seasonal adaptation for multivariate time series forecasting. *Scientific Reports*. 2024;14(1):15909. https://doi.org/10.1038/s41598-024-66886-1

[37] Mejía S, Aguilar J. A demand forecasting system of product categories defined by their time series using a hybrid approach of ensemble learning with feature engineering. *Computing*. 2024;106(12):3945-3965. https://doi.org/10.1007/s00607-024-01320-y

[38] Graham RL, Lawler EL, Lenstra JK, Kan AR. Optimization and approximation in deterministic sequencing and scheduling: a survey. *Annals of Discrete Mathematics*. 1979;5:287-326. https://doi.org/10.1016/s0167-5060(08)70356-x

[39] Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, et al. Human-level control through deep reinforcement learning. *Nature*. 2015;518(7540):529-533. https://doi.org/10.1038/nature14236