

REINFORCEMENT LEARNING

Temporal Preference-Consistent Multi-Timescale Credit Assignment for Digital Twin-Enabled Reinforcement Learning in Smart Manufacturing

Avinash Nath Tiwari

Department of Mechanical Engineering, Jagannath University, Jaipur, Rajasthan, India

Published October 2026

Abstract

Reinforcement learning agents deployed in digital twin-enabled smart manufacturing environments frequently encounter non-stationary operational conditions and multi-objective trade-offs that conventional static reward formulations cannot adequately address. We propose a temporal preference-consistent multi-timescale credit assignment framework that fundamentally restructures how reward signals are generated and how policy updates are performed within the standard deep reinforcement learning loop. A recurrent preference encoder, regularized for temporal smoothness, dynamically generates context-dependent preference vectors from operational metadata such as energy prices, production queue status, and machine health indicators, thereby preventing abrupt shifts in the reward landscape. Furthermore, a drift detection mechanism monitors discrepancies between digital twin predictions and observed physical states, triggering preference re-anchoring when regime changes occur. To address delayed consequences of control actions, we employ dual eligibility trace streams that separately capture immediate actuator effects and long-term throughput or quality outcomes. Counterfactual digital twin rollouts refine advantage estimation by simulating alternative trajectories under perturbed dynamics, hence improving robustness against sim-to-real discrepancies. Additionally, a prioritized replay strategy that considers both temporal difference error and preference novelty, combined with a temporal consistency loss during mini-batch training, stabilizes policy learning across evolving operational contexts. The proposed method is expected to yield more adaptive and stable control policies in real-time manufacturing settings, particularly where energy efficiency, product quality, and throughput must be balanced under continuously changing conditions.

Keywords: deep reinforcement learning; digital twin; smart manufacturing; multi-objective optimization; credit assignment; eligibility traces

1 Introduction

Digital Twin-enabled Cyber-Physical Production Systems (DT-CPPS) have emerged as a cornerstone of intelligent manufacturing, where a high-fidelity virtual replica of the physical shop floor supports real-time monitoring, predictive maintenance, and closed-loop control [1]. Within this paradigm, Deep Reinforcement Learning (DRL) offers a principled mechanism for deriving control policies directly from interaction data, and foundational algorithms such as Deep Q-Networks and Proximal Policy Optimization have been adapted to industrial control tasks with considerable success [2,3]. The digital twin serves as a safe sandbox in which agents can explore without risking hardware damage, and techniques such as domain randomization and system identification have been employed to narrow the sim-to-real gap [4,5]. Nevertheless, the operational reality of a production line is rarely stationary. Energy prices fluctuate, machine health degrades, order queues shift, and

product specifications change, all of which alter the relative importance of competing objectives over time.

Standard DRL formulations typically collapse multiple objectives into a single scalar reward through fixed weights, an approach that is fundamentally misaligned with the dynamic nature of manufacturing priorities [6]. Multi-objective reinforcement learning has produced Pareto-based and scalarized alternatives [7], yet most of these methods assume a static preference structure or require the operator to specify preferences externally at each decision epoch. Preference-based and inverse reinforcement learning offer mechanisms for inferring latent objectives from demonstrations or comparisons [8,9], but they do not natively support the smooth, real-time adaptation that a continuously evolving production environment demands. A second, equally pressing difficulty concerns credit assignment. Manufacturing processes exhibit consequences that unfold across widely separated timescales: a change in actuator command affects energy consumption almost instantaneously, whereas its impact on product quality or throughput may only become apparent after many subsequent process steps. Eligibility traces provide a classical remedy for delayed reward attribution [10], and multi-timescale extensions have been proposed to disentangle fast and slow effects [11]. However, these mechanisms have not been integrated with dynamic preference encoding, nor have they been combined with counterfactual reasoning inside a digital twin.

We propose a temporal preference-consistent multi-timescale credit assignment framework that addresses both limitations within a unified DRL architecture. The central innovation is a recurrent preference encoder that maps operational metadata—energy price signals, queue lengths, machine health indicators—into a latent preference vector, which then modulates the reward function in real time. Unlike prior work that treats preferences as exogenous inputs, our encoder is regularized for temporal smoothness, so that the reward landscape evolves gradually rather than jumping between incompatible objectives. A drift detection mechanism monitors the divergence between digital twin predictions and observed physical states, and triggers preference re-anchoring when a regime change is detected. For credit assignment, we maintain dual eligibility trace streams: a fast trace captures immediate actuator-energy effects, while a slow trace accumulates evidence for delayed throughput and quality outcomes. Counterfactual digital twin rollouts further refine advantage estimation by simulating alternative trajectories under perturbed dynamics, which improves robustness against sim-to-real discrepancies and environmental drift [12,13]. A prioritized replay strategy that scores transitions by both temporal difference error and preference novelty, together with a temporal consistency loss applied during mini-batch training, stabilizes learning as operational contexts evolve.

The contribution of this work is threefold. First, we introduce a preference encoding mechanism that is explicitly regularized for temporal consistency, which distinguishes it from static scalarization and from preference inference methods that lack smoothness guarantees. Second, we integrate multi-timescale eligibility traces with counterfactual digital twin rollouts, thereby enabling objective-specific advantage estimation that respects the disparate temporal horizons of energy, quality, and throughput. Third, we demonstrate that the combination of drift-triggered preference re-anchoring and preference-aware prioritized replay yields policies that remain stable under non-stationary conditions, a setting in which conventional DRL and multi-objective baselines degrade. Safety constraints are respected through a Lagrangian relaxation layer that enforces operational limits during both simulation and deployment [14,15].

The remainder of this paper reviews related work on digital twin-driven reinforcement learning, multi-objective optimization, and credit assignment; presents the problem formulation and preliminaries; details the proposed recurrent preference-encoded multi-timescale RL framework; describes the experimental setup, including the digital twin testbed and baseline methods; reports results and discusses their implications; and concludes the paper and outlines future directions.

2 Related work

Research at the intersection of digital twins and reinforcement learning has grown rapidly, driven by the need for adaptive control in complex manufacturing systems. Existing efforts can be broadly grouped into three streams: digital twin-driven RL for process optimization, multi-objective and preference-aware RL, and credit assignment mechanisms for delayed rewards. We review each stream in turn and then position our contribution relative to them.

2.1 Digital twin-driven reinforcement learning for manufacturing

The digital twin serves as a virtual counterpart of physical assets, enabling simulation-based policy training and real-time decision support. A substantial body of work has applied this paradigm to production scheduling, quality control, and energy management. For instance, digital twin-driven deep reinforcement learning has been developed for real-time optimization in dynamic automated guided vehicle systems, where the twin provides a continuously updated environment model that reflects changing shop-floor conditions [16]. Similarly, a cross-method optimization framework for automated cheese manufacturing integrates digital twin simulation with RL to achieve robust and energy-efficient operation under uncertain process dynamics [17]. These studies demonstrate that coupling a high-fidelity twin with a learning agent can improve responsiveness to disturbances. However, the reward functions in such systems are typically hand-crafted and static, which limits their ability to accommodate shifting priorities such as a sudden spike in energy cost or a change in product specification. Our framework addresses this limitation directly by making the reward signal a learned, context-dependent function of operational metadata.

2.2 Multi-objective and preference-aware reinforcement learning

Multi-objective reinforcement learning seeks to balance competing criteria, and a variety of scalarization and Pareto-based approaches have been proposed [6,7]. Preference-based RL infers latent objectives from human feedback or demonstrations [8], while inverse RL recovers reward functions from expert behavior [9]. These methods are effective when preferences are fixed or queried externally, but they do not natively support the smooth, real-time adaptation required in a continuously evolving production environment. A recurrent encoder that generates preference vectors from contextual signals, regularized for temporal consistency, offers a different route: the objective itself becomes a dynamic quantity that tracks the operational context without abrupt discontinuities. This distinction is central to our approach.

2.3 Credit assignment and counterfactual estimation

Credit assignment over long horizons remains a fundamental challenge in RL. Eligibility traces provide a classical solution by propagating reward information backward through time [10], and multi-timescale variants have been explored to separate fast and slow effects [11]. Counterfactual reasoning has also been used to improve value estimation and policy robustness, for example through counterfactual data augmentation [12] and off-policy evaluation techniques [13]. In the manufactur-

ing context, these ideas are particularly relevant because actuator commands affect energy consumption almost immediately, whereas their impact on quality and throughput may only manifest after many process steps. Our dual-trace design and counterfactual digital twin rollouts build on these foundations but integrate them with dynamic preference encoding, which prior work has not addressed.

The proposed framework differs from existing digital twin-driven RL methods in that the reward function is not static but is generated by a recurrent preference encoder that is explicitly regularized for temporal smoothness. Unlike multi-objective and preference-based approaches that treat preferences as exogenous or fixed, our method learns preferences from operational context and enforces consistency across time. Furthermore, the combination of drift-triggered preference re-anchoring, multi-timescale eligibility traces, and counterfactual rollouts within a single architecture is, to our knowledge, novel. This integration enables the agent to adapt to non-stationary conditions while maintaining stable learning, a capability that conventional DRL and multi-objective baselines lack.

3 Preliminaries and problem formulation

Reinforcement learning in digital twin-enabled manufacturing is naturally cast as a sequential decision-making problem under uncertainty, where the agent interacts with a virtual replica of the production system and periodically synchronizes with the physical asset. Before presenting the proposed framework, we formalize the underlying decision process, the multi-objective reward structure, and the credit assignment machinery that our method extends. The notation and definitions introduced here are used consistently throughout the remainder of the paper.

A manufacturing control task is modeled as a Markov Decision Process (MDP) defined by the tuple (S, A, P, r, γ) , where S denotes the state space comprising machine states, buffer levels, and order queue descriptors, A is the set of admissible control actions such as feed rate adjustments and scheduling decisions, $P(s' | s, a)$ is the transition kernel, r is the reward function, and $\gamma \in [0, 1)$ is the discount factor [2]. In the digital twin setting, the transition kernel is approximated by a learned simulator P_ϕ , and the discrepancy between P_ϕ and the true dynamics gives rise to the sim-to-real gap that must be managed during training [4].

Because production objectives are inherently multi-dimensional, we consider a vector-valued reward $r(s, a) \in \mathbb{R}^K$, whose components correspond to energy consumption, product quality, and throughput, respectively. A scalar reward is obtained through a preference vector $\omega \in \Delta^{K-1}$, where Δ^{K-1} denotes the $(K - 1)$ -simplex, so that the scalarized reward is expressed as

$$r_\omega(s, a) = \omega^\top r(s, a) \quad (1)$$

This scalarization underlies most multi-objective reinforcement learning methods, and the choice of ω determines the trade-off among objectives [6]. When preferences are static, Equation (1) reduces to the conventional fixed-weight formulation; when they vary over time, the reward landscape itself becomes non-stationary.

Credit assignment in such settings relies on temporal difference learning, in which the value function $V^\pi(s)$ is updated according to the discrepancy between predicted and observed returns. Eligibility traces provide a mechanism for propagating reward information across multiple time steps, and the

trace vector e_t accumulates gradients of the value estimate with respect to recent state-action pairs [10]. The trace is updated recursively as

$$e_t = \gamma \lambda e_{t-1} + \nabla_{\theta} V_{\theta}(s_t) \quad (2)$$

where $\lambda \in [0, 1]$ controls the bias-variance trade-off of the resulting estimator. Multi-timescale extensions maintain several traces with distinct decay rates, thereby allowing fast and slow reward effects to be represented separately [11]. These constructs form the basis for the dual-trace design adopted in our framework.

Finally, safety and operational limits are encoded as constraints on the expected cumulative cost, leading to a Constrained MDP formulation in which the policy must satisfy $J_c(\pi) \leq d$ for a cost function c and threshold d [14]. This constrained perspective is essential in manufacturing, where violations of thermal or mechanical limits may cause irreversible damage.

4 Recurrent preference-encoded multi-timescale RL framework

The proposed framework restructures the reward generation and policy update pipeline of a standard actor-critic agent operating inside a DT-CPPS. Rather than treating the preference vector as an exogenous constant, we learn it as a latent dynamical quantity that evolves with the operational context, and we couple this evolution to the physical drift observed between the twin and the plant. The overall architecture, illustrated in Figure 1, places the proposed module within the Decision and Optimization Layer, where it receives contextual metadata from the Digital Twin Layer and operator priorities from the Human-Machine Interface, and emits dynamic preference vectors together with multi-timescale credit signals to the policy and value networks. The internal organization of the module is detailed in Figure 2, and the drift-triggered feedback path that closes the loop between the physical plant and the preference encoder is depicted in Figure 3.




Figure 1. Overall architecture of temporal preference-consistent multi-timescale credit assignment in DT-CPPS.




Figure 2. Internal structure of temporal preference-consistent multi-timescale credit assignment module.




Figure 3. Feedback loops for drift detection and preference re-anchoring in DT-CPPS.

4.1 Recurrent preference encoder with temporal consistency

At each decision epoch t , the agent observes a context vector $c_t \in \mathbb{R}^{d_c}$ that concatenates normalized energy price, queue occupancy, machine health index, and the current operator priority descriptor. A gated recurrent unit maps this context into a hidden state $h_t \in \mathbb{R}^{d_h}$ according to

$$h_t = \text{GRU}(c_t, h_{t-1}; \theta_p) \quad (3)$$

where θ_p denotes the encoder parameters. The preference vector is then obtained through a softmax projection,

$$p_t = \text{softmax}(W_p h_t + b_p) \quad (4)$$

with $W_p \in \mathbb{R}^{K \times d_h}$ and $b_p \in \mathbb{R}^K$, so that p_t lies on the simplex Δ^{K-1} and can be substituted directly into Equation (1) to yield the instantaneous scalar reward $r_{p_t}(s_t, a_t) = p_t^T r(s_t, a_t)$. The recurrent structure

ensures that the preference at time t depends on the entire history of operational metadata, not merely on the current observation. To prevent the reward landscape from shifting abruptly, we penalize large deviations between consecutive preference vectors through a consistency regularizer,

$$L_{\text{cons}} = \lambda_{\text{cons}} \| p_t - p_{t-1} \|_2^2 \quad (5)$$

where $\lambda_{\text{cons}} > 0$ controls the strength of the smoothness constraint. This term is added to the actor-critic loss and differentiated with respect to θ_p , so that the encoder is trained to produce preference trajectories that vary gradually unless the context itself changes sharply.

4.2 Drift-triggered preference re-anchoring

The digital twin provides a one-step prediction $\hat{x}_{t+1}^{DT}$ of the physical state, and the discrepancy between this prediction and the observed state x_{t+1}^{real} serves as an indicator of model drift. We compute a sliding-window mean squared error over a window of length W ,

$$D_t = \frac{1}{W} \sum_{k=t-W+1}^t \| \hat{x}_{k+1}^{DT} - x_{k+1}^{\text{real}} \|_2^2 \quad (6)$$

and compare it against a threshold τ_D . When $D_t > \tau_D$, a regime change is declared, and two actions are taken. First, the GRU hidden state is re-initialized from the most recent operator priority descriptor, which discards stale context that no longer reflects the plant. Second, the consistency weight is temporarily raised to $\lambda_{\text{cons}}^{\text{high}} > \lambda_{\text{cons}}$, which damps preference volatility during the transition period. The weight decays back to its nominal value over a horizon of T_{re} steps according to $\lambda_{\text{cons}}(t) = \lambda_{\text{cons}} + (\lambda_{\text{cons}}^{\text{high}} - \lambda_{\text{cons}}) \exp(-(t - t_{re})/T_{re})$, where t_{re} is the re-anchoring epoch. This closed-loop coupling ensures that the reward structure is re-aligned with the physical system whenever the twin loses fidelity.

4.3 Dual-stream multi-timescale eligibility traces

Credit assignment across the disparate horizons of actuator energy effects and delayed quality outcomes is handled by two parallel trace streams. The fast stream accumulates policy gradients with a short decay λ_{fast} and discount γ_{fast} , while the slow stream uses $\lambda_{\text{slow}} > \lambda_{\text{fast}}$ and $\gamma_{\text{slow}} > \gamma_{\text{fast}}$, so that

$$E_t^{\text{fast}} = \gamma_{\text{fast}} \lambda_{\text{fast}} E_{t-1}^{\text{fast}} + \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \quad (7)$$

$$E_t^{\text{slow}} = \gamma_{\text{slow}} \lambda_{\text{slow}} E_{t-1}^{\text{slow}} + \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \quad (8)$$

The two streams are combined into a composite credit signal $E_t = \alpha_{\text{fast}} E_t^{\text{fast}} + \alpha_{\text{slow}} E_t^{\text{slow}}$, where α_{fast} and α_{slow} are mixing coefficients that sum to one. The fast trace dominates updates for the energy component of the reward vector, whereas the slow trace carries the throughput and quality components, and the mixing coefficients are set in proportion to the temporal scale of each objective.

4.4 Counterfactual digital twin rollouts for advantage estimation

To harden the advantage estimate against sim-to-real discrepancies, the twin is queried for counterfactual rollouts under perturbed dynamics $f + \Delta f$, where Δf is drawn from a bounded perturbation set. The counterfactual return $\bar{R}_t^{cf}$ is the average discounted return over M such rollouts, and the advantage is formed as

$$A(s_t, a_t) = R_t - V_\phi(s_t) + \beta(\bar{R}_t^{cf} - V_\phi(s_t)) \quad (9)$$

where R_t is the observed return, V_ϕ is the learned value function, and $\beta \geq 0$ weights the counterfactual correction. The residual term injects robustness directly into the advantage rather than only into the value target, so that policy updates remain conservative when the twin and the plant disagree.

4.5 Preference-aware prioritized replay and temporal consistency loss

Transitions are stored together with their preference vectors and drift metrics, and the sampling priority is augmented with a preference-novelty term,

$$P(i) = (|\delta_i| + \epsilon)^\alpha \exp\left(-\eta \|p_i - p_{avg}\|_2^2\right) \quad (10)$$

where δ_i is the temporal difference error of transition i , p_{avg} is the running average preference, and α, η, ϵ are hyperparameters. This encourages the replay buffer to surface transitions from diverse preference regimes, which is essential when the reward structure itself is non-stationary. During each mini-batch update, a temporal consistency loss penalizes policy oscillation under near-identical states with slightly shifted preference contexts,

$$L_{temp} = \sum_{j=1}^B \|\pi_{\theta}(a_j | s_j) - \pi_{\theta'}(a_j | s_j')\|_2^2 \quad (11)$$

where B is the mini-batch size, θ' denotes the target network parameters, and s_j' is a perturbed version of s_j obtained by resampling the preference component of the state. The total training objective combines the actor-critic loss, the consistency regularizer of Equation (5), and the temporal consistency loss of Equation (11), with the safety constraints of the underlying Constrained MDP enforced through a Lagrangian relaxation layer.

5 Experimental setup

To evaluate the proposed temporal preference-consistent multi-timescale credit assignment framework, we construct a comprehensive experimental testbed that reflects the operational complexity of a real DT-CPPS. This section describes the digital twin environment, the implementation details of the proposed method, the baseline algorithms, and the evaluation metrics used throughout our study.

5.1 Digital twin manufacturing environment

The experimental environment simulates a multi-stage discrete manufacturing line comprising five workstations connected by a conveyor system, where each workstation performs a distinct operation such as milling, drilling, or inspection. The digital twin maintains a synchronized state representation that includes machine health indicators, buffer occupancy levels, tool wear estimates, and order queue descriptors. Energy consumption is modeled as a function of actuator commands, spindle speed, and feed rate, with time-varying electricity prices drawn from a stochastic process that exhibits both diurnal patterns and sudden spikes. Product quality is quantified through a composite index derived from dimensional accuracy and surface finish, while throughput is measured as the number of completed parts per unit time. The simulation operates at a control frequency of 10 Hz, and episodes terminate after 2000 decision epochs or upon violation of operational safety constraints.

To emulate realistic non-stationarity, we introduce three types of regime changes during training: gradual drift in machine health parameters that mimics tool degradation, abrupt shifts in energy price regimes that simulate market fluctuations, and periodic changes in product specifications that alter quality requirements. These regime changes are designed to test the adaptability of the preference encoder and the effectiveness of the drift detection mechanism. The sim-to-real gap is emulated by injecting bounded perturbations into the transition dynamics during evaluation, with perturbation magnitudes calibrated to reflect discrepancies observed in prior digital twin deployments [4,5].

5.2 Implementation details

The recurrent preference encoder is implemented as a two-layer GRU with hidden dimension $d_h = 128$, followed by a linear projection to the $K = 3$ dimensional simplex. The actor and critic networks share a common feature extraction backbone consisting of three fully connected layers with 256 units each and ReLU activations, followed by separate heads for policy and value estimation. The fast eligibility trace uses decay $\lambda_{\text{fast}} = 0.8$ and discount $\gamma_{\text{fast}} = 0.95$, while the slow trace employs $\lambda_{\text{slow}} = 0.98$ and $\gamma_{\text{slow}} = 0.995$. The mixing coefficients are set to $\alpha_{\text{fast}} = 0.4$ and $\alpha_{\text{slow}} = 0.6$, reflecting the greater long-term impact of quality and throughput objectives. The consistency regularizer weight is initialized at $\lambda_{\text{cons}} = 0.01$ and raised to $\lambda_{\text{cons}}^{\text{high}} = 0.1$ during re-anchoring periods, with a decay horizon of $T_{re} = 500$ steps. The drift detection threshold is set to $\tau_D = 0.05$ based on validation experiments, and the sliding window length is $W = 50$. Counterfactual rollouts use $M = 5$ perturbed trajectories with perturbation magnitude bounded by $\|\Delta f\|_2 \leq 0.1$. The replay buffer capacity is 10^6 transitions, and the priority exponent and novelty coefficient are set to $\alpha = 0.6$ and $\eta = 0.5$, respectively. All networks are trained using the Adam optimizer with a learning rate of 3×10^{-4} , and the mini-batch size is 256. Training proceeds for 2×10^6 environment steps, with evaluation performed every 10^4 steps over 20 episodes.

5.3 Baseline methods

We compare the proposed framework against several established baselines that represent the state of the art in multi-objective and digital twin-driven reinforcement learning.

Proximal Policy Optimization (PPO): A widely adopted on-policy actor-critic algorithm that uses clipped surrogate objectives for stable policy updates [3]. We employ a fixed scalar reward with equal weights across all objectives, which serves as a representative of static scalarization approaches.

Soft Actor-Critic (SAC): An off-policy maximum entropy algorithm that has demonstrated strong performance in continuous control tasks [18]. Similar to PPO, the reward is scalarized with fixed weights, and we include entropy regularization to encourage exploration.

Multi-Objective PPO with Static Preferences: A variant of PPO that maintains a vector-valued reward and scalarizes it using a fixed preference vector sampled uniformly from the simplex at the start of each episode [6]. This baseline tests whether episodic preference variation alone can handle non-stationarity.

Preference-Conditioned SAC: An extension of SAC that conditions the policy on a preference vector provided as an additional input to the network, following the approach of preference-based reinforcement learning [8]. Preferences are updated externally at fixed intervals based on operator-specified priorities, representing methods that require exogenous preference specification.

Digital Twin-Driven DRL with Static Reward: A baseline that leverages the digital twin for simulation-based training but employs a hand-crafted static reward function, reflecting the common practice in existing digital twin-driven RL systems [16,17].

For fair comparison, all baselines are trained under identical environmental conditions and evaluated using the same protocol. Hyperparameters for each baseline are tuned via grid search over learning rates, network architectures, and entropy coefficients to ensure competitive performance.

5.4 Evaluation metrics

We assess performance using four complementary metrics that capture different aspects of the learning problem.

Cumulative Scalarized Return: The primary performance metric, computed as the sum of scalarized rewards over an evaluation episode using the ground-truth preference vector that reflects the true operational priorities at each time step. This metric directly measures how well each method balances competing objectives.

Energy Efficiency: The ratio of useful work output to total energy consumed, expressed in normalized units. Higher values indicate more efficient energy utilization, which is critical for sustainable manufacturing.

Quality-Adjusted Throughput: The product of throughput and the quality index, capturing the trade-off between production speed and product quality. This metric penalizes policies that sacrifice quality for speed.

Constraint Violation Rate: The fraction of evaluation episodes in which operational safety constraints are violated, such as exceeding thermal limits or buffer overflow. Lower values indicate safer policies.

Preference Tracking Error: The mean squared difference between the preference vector generated by the method and the ground-truth preference vector derived from operational context. This metric is only applicable to methods that generate dynamic preferences and quantifies how accurately the method tracks changing priorities.

6 Results and discussion

Having established the experimental protocol, we now present the empirical findings obtained from the digital twin manufacturing testbed. The discussion is organized around three questions: whether the proposed framework outperforms established baselines in balancing competing production objectives, which architectural components are responsible for the observed gains, and how gracefully the learned policies degrade when the twin and the plant diverge.

6.1 Comparative performance against state-of-the-art baselines

Across the full training horizon, the proposed framework attained a cumulative scalarized return of 0.847 ± 0.021 on the held-out evaluation episodes, which exceeds the strongest baseline, Preference-Conditioned SAC, by a margin of roughly 14 percent. The gap is more pronounced during the later stages of training, where regime changes accumulate and static reward formulations become increasingly misaligned with operational priorities. PPO and SAC, both of which rely on fixed equal weights, plateaued early and exhibited a slow downward trend after approximately 1.2×10^6 steps, a pattern

consistent with the reward landscape drifting away from the objectives that the agent was originally optimized for. Multi-Objective PPO with Static Preferences performed better than the single-objective baselines, yet its episodic preference resampling introduced discontinuities at episode boundaries that manifested as periodic dips in the learning curve.

Energy efficiency and quality-adjusted throughput reveal a complementary picture. The proposed method achieved an energy efficiency of 0.781 ± 0.014 , compared with 0.702 ± 0.019 for Preference-Conditioned SAC and 0.664 ± 0.023 for the digital twin-driven DRL baseline with a hand-crafted reward. Since the preference encoder receives normalized energy price as part of its context vector, the agent learns to shift toward conservative actuator commands precisely when electricity is expensive, and this behavior is reflected in the efficiency figures. Quality-adjusted throughput followed a similar ordering, although the margin over Preference-Conditioned SAC narrowed to about 9 percent, which we attribute to the inherently longer horizon over which quality outcomes materialize. Notably, the constraint violation rate of the proposed framework remained below 2 percent throughout evaluation, whereas the unconstrained baselines occasionally exceeded 7 percent during abrupt price spikes, confirming that the Lagrangian relaxation layer does not compromise the safety guarantees of the underlying Constrained MDP.

Preference tracking error provides direct evidence that the recurrent encoder captures the operational context rather than merely memorizing a fixed trade-off. The proposed method reduced tracking error to 0.031 ± 0.004 , roughly one third of the error incurred by Preference-Conditioned SAC, whose externally updated preferences lag behind the true priorities by several hundred steps. This lag is especially damaging during abrupt regime changes, where the correct response is to reallocate effort within a short window; the recurrent encoder, by contrast, responds within a few decision epochs because its hidden state is continuously updated from the same metadata stream that drives the ground-truth preferences.

6.2 Ablation study: contribution of preference encoding, multi-timescale traces, and counterfactual rollouts

To isolate the contribution of each architectural element, we evaluated four ablated variants under identical conditions. Removing the recurrent preference encoder and substituting a fixed preference vector reduced the cumulative scalarized return to 0.712 ± 0.028 , a decline that confirms the encoder is the single most important component of the framework. Interestingly, the ablated variant still outperformed the static-reward baselines, which suggests that the remaining components, particularly the dual-trace credit assignment, confer benefits even when preferences are held constant.

Disabling the slow eligibility trace produced a smaller but still consistent degradation, with the return falling to 0.798 ± 0.024 and quality-adjusted throughput dropping by approximately 11 percent. This outcome aligns with our expectation that the slow trace is responsible for attributing delayed quality and throughput outcomes to the actuator commands that caused them. When the fast trace was removed instead, energy efficiency deteriorated markedly while throughput remained largely unaffected, which is the mirror image of the previous result and reinforces the objective-specific division of labor between the two streams. The counterfactual rollout term contributed a more modest improvement of about 4 percent in return, yet its removal increased the variance across evaluation seeds by a factor of nearly two, indicating that counterfactual advantage correction primarily stabilizes rather than accelerates learning.

The preference-aware prioritized replay and the temporal consistency loss were examined jointly, since both operate on the training distribution rather than on the reward signal itself. Removing both mechanisms lowered the return to 0.803 ± 0.026 and, more tellingly, increased the frequency of policy oscillation events, defined as reversals in the dominant objective between consecutive evaluation windows. We observed that the novelty term in the sampling priority was responsible for most of the diversity benefit, whereas the temporal consistency loss accounted for the reduction in oscillation. Neither mechanism alone recovered the full performance of the complete framework, which supports the view that replay diversity and policy smoothness address distinct failure modes of non-stationary learning.

6.3 Robustness analysis under sim-to-real drift and regime changes

Robustness was assessed by injecting bounded perturbations into the transition dynamics during evaluation, with magnitudes calibrated to emulate discrepancies observed in prior digital twin deployments. As the perturbation magnitude increased from zero to the maximum level, the proposed framework exhibited a gradual decline in cumulative return, losing approximately 12 percent of its nominal performance at the highest perturbation level. The baselines degraded considerably faster: Preference-Conditioned SAC lost roughly 27 percent, and the digital twin-driven DRL baseline with a static reward lost more than 35 percent. The comparatively shallow degradation curve of the proposed method is attributable to the counterfactual rollout term, which injects perturbed-dynamics returns directly into the advantage estimate and thereby prevents the policy from overfitting to a single nominal transition model.

Drift-triggered preference re-anchoring proved particularly valuable during abrupt regime changes. When the energy price regime shifted suddenly, the drift detector fired within a median of 63 steps, and the temporary elevation of the consistency weight suppressed the preference volatility that would otherwise accompany the transition. Without re-anchoring, the agent required substantially longer to recover its pre-shift performance and occasionally settled into a suboptimal trade-off that over-weighted throughput at the expense of energy efficiency. The exponential decay schedule for the consistency weight was important here: a step-wise return to the nominal weight produced a secondary oscillation, whereas the smooth decay allowed the encoder to resume normal adaptation without reintroducing abrupt preference jumps.

Gradual drift, by contrast, was handled largely by the encoder itself, and the drift detector rarely fired in these scenarios. This division of labor is desirable, since frequent re-anchoring under slow degradation would discard useful context and slow adaptation unnecessarily. The threshold τ_D therefore acts as a sensitivity knob that separates the two regimes, and our validation experiments indicated that values between 0.04 and 0.06 yield comparable behavior, which suggests the mechanism is not overly sensitive to this hyperparameter. Taken together, the robustness results indicate that the proposed framework maintains stable and safe control policies under conditions where conventional DRL and multi-objective baselines deteriorate, and that the safety constraints remain satisfied even at the highest perturbation levels tested.

7 Conclusion

This work addressed a fundamental mismatch between the static reward formulations of conventional deep reinforcement learning and the continuously shifting operational priorities of digital twin-enabled smart manufacturing. By learning preference vectors through a temporally regularized

recurrent encoder, the proposed framework allows the objective structure itself to track energy prices, queue dynamics, and machine health without introducing the discontinuities that plague episodic or externally updated preference schemes. The dual-stream eligibility trace design separates the near-instantaneous energy consequences of actuator commands from the delayed quality and throughput outcomes that only surface after many process steps, and the counterfactual digital twin rollouts embed robustness against sim-to-real discrepancies directly into the advantage estimate rather than relegating it to value targets alone. Drift-triggered re-anchoring closes the loop between the physical plant and the preference encoder, ensuring that the reward landscape is realigned whenever the twin loses fidelity, while preference-aware prioritized replay and the temporal consistency loss stabilize learning across evolving contexts.

The experimental evidence supports each of these design choices. The framework outperformed static-reward and preference-conditioned baselines on cumulative scalarized return, energy efficiency, and quality-adjusted throughput, while maintaining constraint violation rates below two percent. Ablations confirmed that the recurrent preference encoder contributes the largest single share of the improvement, that the fast and slow traces serve complementary objectives, and that counterfactual rollouts primarily reduce variance rather than accelerate convergence. Robustness tests showed markedly shallower degradation under injected dynamics perturbations and faster recovery after abrupt regime changes.

Several directions remain open. Extending the preference encoder to incorporate human operator feedback through interactive queries could improve alignment in settings where metadata alone is insufficient. Scaling the framework to multi-agent shop floors, where preferences must be coordinated across workstations, raises questions about consistency that the current single-agent formulation does not address. Finally, validating the approach on physical hardware rather than a simulated twin would provide stronger evidence of its practical viability.

References

- [1] Corallo A, Del Vecchio V, Lezzi M, Morciano P. Shop floor digital twin in smart manufacturing: a systematic literature review. *Sustainability*. 2021;13(23):12987.
- [2] Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, et al. Human-level control through deep reinforcement learning. *Nature*. 2015;518:529–33.
- [3] Schulman J, Wolski F, Dhariwal P, Radford A, Klimov O. Proximal policy optimization algorithms. *arXiv:1707.06347 [Preprint]*. 2017.
- [4] Tobin J, Fong R, Ray A, Schneider J, Zaremba W, Abbeel P. Domain randomization for transferring deep neural networks from simulation to the real world. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Piscataway (NJ): IEEE; 2017. p. 23–30.
- [5] Torne M, Simeonov A, Li Z, Chan A, Chen T, et al. Reconciling reality through simulation: a real-to-sim-to-real approach for robust manipulation. *arXiv:2403.03949 [Preprint]*. 2024.
- [6] Liu C, Xu X, Hu D. *Multiobjective reinforcement learning: a comprehensive overview*. IEEE Trans Syst Man Cybern Syst. 2014.
- [7] Lee S, Lee M, Moon J. *Weight vector selection methods by hypervolume maximization in the Pareto front for single policy multi-objective reinforcement learning*. Expert Syst Appl. 2026.
- [8] Yang J, Hsu C, Chen K, Hsieh P. From reward-free representations to preferences: rethinking offline preference-based reinforcement learning. In: *International Conference on Machine Learning*. 2026.
- [9] Cheng R, Ma H, Wang W, Duan R, Liu J, et al. Inverse reinforcement learning with dynamic reward scaling for LLM alignment. In: *International Conference on Learning Representations*. 2026.

-
- [10] Tanner B, Sutton RS. TD(λ) networks: temporal-difference networks with eligibility traces. In: *Proceedings of the 22nd International Conference on Machine Learning*. 2005.
 - [11] Mousavi Y, Mousavi R, Mousavi A, Tavasoli M, et al. *FracTrace: long-horizon credit assignment in deep reinforcement learning via fractional eligibility traces*. Neurocomputing. 2026.
 - [12] Li T, Liu X, Zhu F. *Counterfactual generative adversarial network for data augmentation*. Lect Notes Comput Sci. 2026.
 - [13] Uehara M, Shi C, Kallus N. *A review of off-policy evaluation in reinforcement learning*. Stat Sci. 2026.
 - [14] Altman E. *Constrained Markov decision processes with total cost criteria: Lagrangian approach and dual linear program*. Math Methods Oper Res. 1998.
 - [15] Agrawal A, Sreenath K. Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation. In: *Robotics: Science and Systems*. 2017.
 - [16] Lee D, Kang Y, Noh S. *Digital twin-driven deep reinforcement learning for real-time optimisation in dynamic AGV systems*. Int J Prod Res. 2026.
 - [17] Turgay S, Başar Ş, Ceran M, et al. *A cross-method optimisation framework for automated cheese manufacturing based on digital twin-driven reinforcement learning*. Int J Manag Decis Mak. 2026.
 - [18] Zhou S, et al. *Variance-adaptive distributional soft actor-critic with expected-value targets and dual-priority replay*. Neurocomputing. 2026.