

Logic-informed reinforcement learning for cross-domain optimization of large-scale cyber-physical systems

Guangxi Wan¹, Peng Zeng^{1,*}, Xiaoting Dong^{1,2}, Chunhe Song¹, Shijie Cui¹, Dong Li¹, Qingwei Dong¹, Yiyang Liu¹, and Hongfei Bai¹

¹State Key Laboratory of Robotics and Intelligent Systems, Shenyang Institute of Automation, Chinese Academy of Sciences, Shenyang 110016, China

²University of Chinese Academy of Sciences, Beijing 100049, China

*Corresponding: zp@sia.cn

ABSTRACT

Cyber-physical systems (CPS) require the joint optimization of discrete cyber actions and continuous physical parameters under stringent safety logic constraints. However, existing hierarchical approaches often compromise global optimality, whereas reinforcement learning (RL) in hybrid action spaces often relies on brittle reward penalties, masking, or shielding and struggles to guarantee constraint satisfaction. We present logic-informed reinforcement learning (LIRL), which equips standard policy-gradient algorithms with projection that maps a low-dimensional latent action onto the admissible hybrid manifold defined on-the-fly by first-order logic. This guarantees feasibility of every exploratory step without penalty tuning. Experimental evaluations have been conducted across multiple scenarios, including industrial manufacturing, electric vehicle charging stations, and traffic signal control, in all of which the proposed method outperforms existing hierarchical optimization approaches. Taking a robotic reducer assembly system in industrial manufacturing as an example, LIRL achieves a 36.47%–44.33% reduction at most in the combined makespan–energy objective compared to conventional industrial hierarchical scheduling methods. Meanwhile, it consistently maintains zero constraint violations and significantly surpasses state-of-the-art hybrid-action reinforcement learning baselines. Thanks to its declarative logic-based constraint formulation, the framework can be seamlessly transferred to other domains such as smart transportation and smart grid, thereby paving the way for safe and real-time optimization in large-scale CPS.

Cyber-physical systems (CPS) seamlessly integrate sensing, computation, and actuation to enable closed-loop control of large-scale physical processes^{1–4}. Prominent examples, such as smart factories⁵, autonomous transportation networks⁶, and wide-area power grids⁷, constitute crucial components of modern industry and infrastructure. In these applications, the overall performance of the system is critically based on cross-domain optimization⁸. At each decision point, the supervisory controller must simultaneously (i) schedule discrete cyber issues, including task assignments, mode switches, and component start-ups, and (ii) adjust continuous physical parameters, such as robot trajectories, vehicle accelerations, or generator set-points. These two aspects are interlinked by strict safety, logical, and resource constraints. Violations can result in severe system-wide failures. Therefore, attaining real-time and globally optimal decision-making within such hybrid decision spaces poses a fundamental challenge in both science and engineering.

In contemporary industrial practices, a significant reliance is placed on hierarchical decoupling⁹. Specifically, a high-level planner first devises a discrete schedule and then conveys it to a low-level optimizer, or the process may occur in the reverse order. This low-level optimizer refines continuous parameters under the premise of fixed logic. Although this approach is computationally manageable, it unavoidably forfeits global optimality, as shown in Figure 1A. This is primarily because the optimization of each layer is often based on idealized assumptions, failing to adequately account for mutual coupling constraints and dynamic interactions. Moreover, no optimal substructure exists between the cyber and physical layer optimizations in the decomposed sub-problems. As manufacturing, traffic, and energy systems continue to expand on a scale, the performance degradation that stems from decoupling becomes increasingly prominent¹⁰.

Reinforcement learning (RL) presents an appealing alternative as, in theory, it can learn direct mappings from system state to hybrid actions without explicitly modeling the underlying dynamics. RL has achieved superhuman performance in various domains, including games, robotic manipulation, and the energy management of single electric vehicles. However, its application to large-scale CPS has been hampered by two persistent challenges. Firstly, the decision space in such scenarios is hybrid and high-dimensional. The discrete and continuous components demand distinct policy representations, and the size




Figure 1. Overview of CPS cross-domain optimization and LIRL framework. **A.** The concept of Hierarchical optimization and cross-domain optimization. Hierarchical optimization struggles to achieve the global optimum due to the lack of a strict optimal substructure between levels, while cross-domain optimization can attain the globally optimal solution. **B.** The LIRL introduces two key innovations: latent action projection and dynamic valid-action-space partitioning. It first maps hybrid discrete-continuous actions into a unified latent continuous space, then projects them via function Π_s into the valid explicit action space—which is adaptively partitioned based on system state and constraints. **C.** The characteristics of the LIRL training process guarantee the continuous satisfaction of constraints throughout the training period, while steadily enhancing performance. **D.** The process of latent action projection and valid action space partitioning. The valid action space represents the constrained solution space under varying CPS system states. Decision-making activates only when the valid action space is non-empty and halts automatically otherwise.

of their Cartesian product expands exponentially with the growth of the system. Specialized architectures like parameterized deep Q-Network¹¹ (PDQN), hierarchical proximal policy optimization¹² (HPPO), or multi-branch actor-critic networks can address specific subsets of hybrid problems. However, these architectures suffer from limited generality and complex network design requirements. Secondly, RL depends on reward signals to drive exploration. In CPS, objectives, such as minimizing

a weighted combination of energy consumption, throughput, and safety risks, are intricately linked with strict logical and physical constraints. Encoding these constraints into a single scalar reward requires a delicate balance between positive goals and negative penalties. Incorrect specification often results in either continuous constraint violations or overly cautious policies that sacrifice performance. Although augmenting RL with action masking or shielding can alleviate immediate violations, it introduces discontinuous gradients, which impede convergence and may lead to short-sighted corrective actions.

Recent investigations in the realm of logic-guided or safe RL have incorporated temporal-logic templates or safety layers into policy learning, thus ensuring the satisfaction of constraints^{13,14}. Nevertheless, the majority of these approaches are tailored either to purely discrete actions^{15,16} or to purely continuous control with differentiable constraints^{17–19}. They do not offer a unified approach for hybrid CPS and do not analyze how the imposed corrections interact with gradient-based policy updates.

In light of these gaps, we propose logic-informed reinforcement learning (LIRL), a general framework shown in Figure 1B that enhances off-the-shelf policy-gradient algorithms to optimize large-scale CPS under strict adherence to hybrid constraints. The core idea is to decouple exploration from feasibility through a latent unified action space. At each decision-making step, the agent proposes a low-dimensional latent vector, which is then projected onto the admissible manifold defined by first-order logic incorporating both cyber and physical constraints. This projection operation (i) guarantees that all executed actions are feasible, (ii) preserves gradient smoothness in continuous sub-spaces, and (iii) eliminates the need for reward shaping or pre-trained autoencoders. This paradigm ensures continuous satisfaction of constraints throughout training, while steadily improving performance, as shown in Figure 1C. Unlike traditional RL methods, where exploratory actions often cause early constraint violations, our approach guarantees strict constraint compliance from the very beginning, even when using random policies. As Figure 1D illustrates, we define a state-specific valid action space. The presence of valid actions determines the decision point: optimization is triggered when available; otherwise, the cycle terminates. This mechanism restricts the agent to only generate valid actions, thereby concentrating exploration entirely on feasible policies. By updating the policy solely with valid actions, the policy gradient is guided toward high-reward and constraint-satisfying regions, accelerating convergence. Since the logical specification is declarative, the same reinforcement learning backbone can be efficiently applied to various CPS domains with minimal engineering effort.

Our approach delivers four main contributions in this paper: (i) The cross-domain optimization of CPS is formulated as a constrained Markov decision process²⁰ with hybrid actions, and a projection-based method is presented. This method transforms the latent action into a hybrid action that satisfies the constraints in a closed form. (ii) Under mild Lipschitz conditions, the projection operator is proven to maintain the fixed points of standard policy-gradient updates, thereby preserving the convergence guarantees of the underlying RL algorithm. (iii) LIRL was implemented on a robotic reducer-assembly manufacturing system and evaluated in four optimization scenarios designed to jointly minimize makespan and energy consumption. Compared with state-of-the-art hierarchical schedulers, LIRL achieves an improvement of 36.47%–44.33% at most in aggregate reward. It consistently outperforms recent hybrid-action RL baselines and completely avoids constraint violations throughout the entire training process. (iv) Extensive ablation studies and cross-domain case studies—including elevator door header assembly, traffic signal coordination, and electric vehicle (EV) charging station scheduling—were systematically conducted to validate the generality, sample efficiency, and scalability of LIRL.

Results

Empirical evidence from the simulation of a robotic reducer-assembly manufacturing System (R^2AMS) (Figure 2) demonstrates the effectiveness, efficiency, and robustness of the LIRL approach. All code, configuration files, and raw logs are publicly accessible at <https://github.com/wanguangxi/LIRL-CPS>. Details of additional experiments carried out in a real elevator door assembly industry, a simulated transportation scenario, and a simulated micro-grid scenario are provided in the Supplementary Material.

Overview of the robotic reducer-assembly system

Gear assembly accounts for 35% of all industrial assembly tasks²¹, representing a canonical case in manufacturing systems²². We therefore selected the reducer—one of the most representative products in gear-based assembly—as our experimental testbed. The R^2AMS consists of a set of robot workcells, as shown in Figure 2A, which is tasked with assembling a batch of reducer products. To further evaluate the generalization capability of the LIRL policy, we designed two dynamic scenarios: one introduces uncertainty in robot operation time due to noise, and the other involves robot failures that lead to uncertainty in production resource availability. The optimization objective of R^2AMS is formulated as a multi-objective function targeting the minimization of both makespan and robotic energy consumption, which is expressed as a weighted combination of these two criteria. An individual reducer assembly task encompasses five stages: *place bottom shell*, *small gear assembling*, *large gear assembling*, *top bottom shell* and *move reducer to buffer*, as shown in Figure 2B. Each robot workcell is capable of performing a complete reducer assembly task. Consequently, R^2AMS represents a fully flexible production and manufacturing system.




Figure 2. Overview of R^2AMS . **A.** R^2AMS consists of modular robotic workcells, each considered a subsystem. To evaluate the generalization capability of LIRL, we designed two dynamic scenarios: task duration uncertainty and robot breakdown. **B.** Five steps for reducer assembly: *place bottom shell*, *small gear assembling*, *large gear assembling*, *top bottom shell* and *move reducer to buffer*. **C.** The structure of hybrid decision-making space. T_i is the i -th reducer; p_{ij} is the j -th stage of the i -th reducer assembly task; q_k is the k -th robot workcell; u_k^{ij} represents the trajectory configuration parameters of the robot workcell q_k when completing p_{ij} . **D.** The energy signature of a robot workcell to process five operations of reducer assembling. The horizontal axis corresponds to the action processing time, where shorter times correspond to faster robot execution.

Decision variables

The decision-making space for R^2AMS cross-domain optimization is a hybrid discrete-continuous space, as shown in Figure 2C. The hybrid action $a = (a_d, a_c)$ contains a discrete tuple $a_d = \langle p_{ij}, q_k \rangle$ that assigns a reducer job to a workcell in every operation and a continuous vector $a_c = u_k^{ij}$ that specifies the spline interpolation parameters of the robot trajectory (Methods). The structure of the solution space T_i represents the i -th reducer, p_{ij} represents the j -th operation of the i -th reducer, and u_k^{ij} represents the trajectory parameters of the robot workcell q_k when completed p_{ij} . The logical specification encodes (i) capacity constraints, (ii) precedence constraints across the five operations, and (iii) joint-velocity-limit (Kinematics) constraints for trajectories.

The execution trajectory of a robot workcell significantly influences both operational processing time and energy consumption. Simulation experiments quantified the energy consumption profile of a single robot workcell across operations (Figure 2D). The results demonstrate a non-linear relationship between energy consumption and task completion time for a specific operation: At shorter durations, high-speed motion with frequent acceleration/deceleration elevates energy consumption; as completion time increases, reduced acceleration/deceleration minimizes non-functional losses through smoother joint trajectories; however, at excessively prolonged durations, sustained posture maintenance increases non-functional energy consumption, leading to renewed energy escalation. The nonlinear energy-time relationship in robotic operations underscores the need for cross-domain

optimization of CPS.

Optimization objectives

We minimize a convex combination $R = -(\alpha \times C_{makespan} + (1 - \alpha) \times E_{total})$ where $\alpha \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$. Makespan is measured in seconds and energy in Joule, both normalized to zero mean and unit variance before weighting.

Baselines

We compare LIRL with seven baselines representative of current practices:

- *Energy-opt* and *Time-opt*: Hierarchical schedulers that first fix the robot velocities to energy-optimal or time-optimal values and then solve mixed-integer linear programs for task allocation.
- HyAR²³: Latent-space reinforcement learning with conditional variational autoencoder (CVAE) pre-training.
- PDQN and HPPO: Specialized reinforcement learning algorithms for hybrid-action spaces.
- SAC-Lag²⁴: Soft Actor-Critic with Lagrangian constraint relaxation.
- CPO²⁵: Constrained policy optimization enforcing safety constraints directly in policy space via second-order approximation.

All RL methods are trained for 1000 environment steps and are averaged over 10 random seeds.

Simulation Results: overall performance

The final performance across four production scales, namely *J10R3*, *J20R3*, *J50R5*, and *J100R5*, with different weightings is illustrated in Figure 3A. Here, *J100R5* represents 100 reducer assembly tasks carried out in 5 robotic workcells.

The empirical results reported in Figure 3A and Figure 3B can be synthesized into five principal findings, which collectively corroborate the effectiveness and scalability of the proposed *Cross-opt* (LIRL) strategy for multi-scale flexible assembly systems.

Across all four production scales (*J10R3*, *J20R3*, *J50R5*, and *J100R5*) and for the entire weighting range $\alpha \in [0.1, 0.9]$, *Cross-opt* consistently attains the highest cumulative reward. Its confidence intervals fully envelope those of *Energy-opt* and *Time-opt*, and the linear-regression slopes of its reward curves are the smallest. These observations indicate that *Cross-opt* persistently approaches the global optimum of the feasible space for the bi-criteria objective (makespan versus energy), whereas the baselines converge to sub-optimal solutions.


The rewards of *Energy-opt* and *Time-opt* decrease monotonically with increasing α , and the decline steepens markedly for larger problem instances. In sharp contrast, *Cross-opt* exhibits an almost flat reward curve. For the *J20R3* instance, when α is varied from 0.1 to 0.9, *Cross-opt* fluctuates by merely 4.67%, whereas *Energy-opt* and *Time-opt* deteriorate by 54.21% and 62.38%, respectively. Hence, *Cross-opt* can operate stably in workshops where the preference weights are uncertain or dynamically adjusted.

After normalizing the composite objective $L = \alpha C_{makespan} + (1 - \alpha) E_{total}$, we observe the following: On the medium-scale *J20R3* case, *Cross-opt* achieved performance improvements of 14.89% and 10.11% over *Energy-opt* and *Time-opt*, respectively, at $\alpha = 0.5$. The advantage widens to 44.33% and 36.47% at $\alpha = 0.9$. On the large-scale *J100R5* case, the same tendency persists and does not attenuate with the increased number of tasks and robots. These evidences confirm that *Cross-opt* can be transferred across scales with (approximately) linear growth in computational effort.

The baseline methods search solely in the *cyber layer* (job-station assignment and sequencing), whereas the *physical layer* (robot trajectory and dynamics) remains fixed. Consequently, the degrees of freedom are insufficient when α changes, leading to a rapid degradation of their Pareto frontiers. *Cross-opt*, empowered by reinforcement learning, jointly explores both layers, thereby performing high-dimensional, cooperative search that continuously reshapes the Pareto frontier and preserves the performance upper bound.

For manufacturing scenarios with uncertain or frequently re-tuned preference weights, *CROSS-OPT* can substantially reduce model maintenance and parameter calibration costs. As the scale of flexible manufacturing expands, traditional single-layer or serial optimization frameworks become susceptible to the *curse of dimensionality*, whereas the LIRL paradigm offers a viable, scalable alternative.

In summary, *Cross-opt* outperforms *Energy-opt* and *Time-opt* in terms of optimality, robustness, and scalability, thereby substantiating the potential of cross-domain optimization for complex manufacturing systems.



Convergence and sample efficiency

Because every production-scheduling constraint is hard by definition, designing a well-behaved reward function is far from trivial. In practice, the original baseline agents could not eliminate constraint violations and therefore failed to converge. To guarantee feasibility at every decision step, each baseline was augmented with an invalid action mask²⁶ layer, yielding the variant referred to as baseline+Mask, such as HyAR+Mask in Fig. 3C. These masked baselines were then benchmarked against the proposed LIRL, under *J100R5* with $\alpha = 0.5$.

The quantitative results reported in Figure 3C.I, Figure 3C.II, and Figure 3C.III give rise to five main observations.

Throughout training LIRL attains the largest average return among all competitors and culminates in the highest final reward (Figure 3C.I). Given that all algorithms are subject to identical safety constraints, these results indicate that the policy returned by LIRL is closest to the global optimum within the feasible action space.

Kernel Density Estimation of the terminal returns (Figure 3C.II) yields a standard deviation of $\sigma = 0.8963$ for LIRL, which is the smallest among all baselines. The tight concentration across 10 random seeds implies superior run-to-run consistency and, by extension, higher robustness in real-world deployments.

As shown in Figure 3C.III, LIRL requires only 22 episodes to reach 95% of its final performance, whereas the next fastest method, PDQN+Mask, needs 181 episodes. SAC-Lag+Mask and CPO+Mask further lag behind at 656 and 948 episodes, respectively. Hence, LIRL accelerates convergence by an order of magnitude (10–40 $\times$), primarily because exploration is confined to the constraint-satisfying sub-space, avoiding low-reward actions that violate safety.

HyAR+Mask exhibits a smooth early climb but relies on a CVAE pre-training stage that already exploits LIRL trajectories, yielding a near-optimal starting point; it is nevertheless surpassed by LIRL in later stages. Penalty-based methods (SAC-

Lag+Mask, CPO+Mask) must repeatedly sample infeasible actions and dynamically adjust Lagrange multipliers, leading to slow and oscillatory learning curves. Although PDQN and HPPO profit from masking, they still depend on ϵ -greedy or Gaussian exploration, resulting in larger terminal variance ($\sigma \approx 4\text{--}22$).

Explicitly encoding safety as an action-generation prior while retaining a reward-maximization objective yields a favorable triple outcome: high efficiency, high return, and high stability. In domains where interaction is costly or failure is unacceptable (e.g. robotic manipulation or autonomous driving), the order-of-magnitude speed-up achieved by LIRL translates directly into reduced tuning costs and enhanced operational reliability.

In summary, LIRL simultaneously dominates strong baselines on optimality, stability, and convergence speed, thus furnishing a promising solution for high-dimensional, safety-constrained control.

Ablation study

This study employs the deep deterministic policy gradient²⁷ (DDPG) framework to compare two action-processing strategies under an identical agent policy. The first strategy applies a projection method grounded in LIRL, whereas the second adopts an invalid-action masking mechanism. Experiments are conducted on four task configurations—*J10R3*, *J20R3*, *J50R5*, and *J100R5*—with the weight parameter α fixed at 0.5.




Figure 4. Ablation results. **A.** The learning curve comparison of four scales. **B.** The post-convergence reward distribution. **C.** Comparison of the average rewards of LIRL and Mask under different scales

The learning curves in Figure 4A reveal a markedly steeper ascent for LIRL during the first 150 episodes, after which the return quickly plateaus. Conversely, the Mask baseline exhibits a step-wise, slow climb and has not fully stabilized by the end of the 1 000-episode horizon. The gap is particularly salient in the high-dimensional tasks (50×5 and 100×5), underscoring the superior sample efficiency of LIRL.

The final-episode average returns (Figure 4C) demonstrate a systematic advantage of LIRL that widens monotonically with task scale: 10×3 : -9.9 vs. -10.5 ; 20×3 : -15.3 vs. -16.9 ; 50×5 : -20.5 vs. -26.7 ; 100×5 : -31.9 vs. -67.2 . The reward gap therefore expands from 0.6 to 35.3, indicating that LIRL scales more gracefully to exponentially larger state-action spaces.

Kernel density estimates of the return distribution (Figure 4B) show that LIRL yields a higher and sharper peak, while its tails are substantially curtailed; the empirical standard deviation is roughly 50% that of the Mask baseline. Hence, LIRL not only improves the expected performance but also reduces variance, evidencing greater training stability.

The projection confines policy outputs to the convex set of legal actions, thereby eliminating the zero-gradient regions introduced by masking and retaining continuous, informative gradients—accounting for the accelerated convergence. By suppressing exploration of infeasible actions at the source, LIRL increases the density of informative samples, translating into higher final returns under identical interaction budgets. As the action space grows, Mask still performs forward passes over all actions before zeroing invalid ones, incurring redundant computation and additional noise. LIRL mitigates both issues, yielding super-linear relative gains in large-scale scenarios.




Figure 5. Results of the robustness testing experiment. The "Training" performance refers to models trained and evaluated at specific perturbation levels, whereas the "Generalization" performance is assessed on a single model trained under disturbance-free conditions. "Cov" indicates that the solution obtained through "Generalization" covers the range of the "Training" solution. **A.** Performance of LIRL-trained strategies under different noise levels. **B.** Performance of LIRL-trained strategies under different robot failure rate.

Across convergence speed, asymptotic return and stability, LIRL delivers statistically significant improvements over Invalid-Action Masking, with benefits that amplify as task complexity increases. These findings substantiate the suitability of LIRL for large-scale constrained reinforcement learning problems.

Robustness to stochastic disturbances

Real-world manufacturing lines suffer from uncertain processing times and unexpected machine faults. We therefore introduce stochastic noise perturbation and exogenous failure into simulation.

- **Duration time uncertainty.** The processing time of each operation is perturbed by zero-mean Gaussian noise with standard deviations set to 0.1μ , 0.3μ and 0.5μ , where μ represents the average completion time of each operation, $\mu = t_{min} + 0.5 \times (t_{max} - t_{min})$, t_{min} , t_{max} denote the minimum and maximum theoretical duration times of each operation, respectively.
- **Robot breakdown.** Robots are assumed to break down with probabilities of 1 %, 3 %, and 5 % per robot workcell. Each failure triggers a repair period lasting between $1.5 \times$ and $3 \times$ the nominal processing time of the interrupted task. At most ten failures are allowed within a single training episode.

Figure 5 shows the complete performance distributions obtained with the proposed LIRL policy under two typical out-of-distribution (OOD) perturbations: additive input noise and stochastic machine failures. The curves marked *Training* denote in-distribution performance, whereas *Generalization* reports the same policy evaluated on perturbed test instances. Across noise levels $\mu = 0.1 - 0.5$ and failure rates 1%–5%, the *Generalization* curves cover more than 90% of the *Training* range; for $\mu = 0.5$ and a 3% failure rate the coverage even reaches 100%. For all three metrics—*Makespan*, *Total Energy* and *Weighted Performance*—the discrepancy between the two means never exceeds 0.075, implying that, under medium-to-high OOD perturbations, the policy remains well confined to the training distribution with only moderate variance.




Figure 6. Overview and result of validation in an actual elevator door head manufacturing factory. **A.** Overview of the elevator door-header plant. The factory contains two parallel assembly lines, and each line is equipped with an identical hanging-plate assembly workcell. Every workcell comprises two industrial robots and one dedicated assembly machine. The robots collaborate to execute three sequential operations for each workpiece: (O_1) tension-spring-bracket assembly, (O_2) door-lock assembly, and (O_3) hanging-plate transfer. **B.** The original scheduling plan of the factory. **C.** The optimized scheduling plan generated by LIRL.

As perturbations intensify, both *Training* and *Generalization* distributions shift rightward simultaneously, while the pairwise distance between them remains nearly constant. In particular, the energy metric attains 100% coverage at all three failure levels, demonstrating that the learned policy can stabilize energy consumption without compromising production tempo.

When $\mu = 0.5$ or the failure rate rises to 5%, the mean gap increases once to about 0.07–0.08; nevertheless, the coverage is still no less than 90%. Hence the policy does not collapse but rather enters an acceptable-degradation zone. A single LIRL policy trained on clean data can be deployed directly over a wide spectrum of perturbation magnitudes, eliminating the need for multiple models and repeated hyper-parameter tuning in industrial settings. Even under sensor drift, equipment ageing, or other unforeseen factors, the policy sustains near-optimal cycle time and energy usage, thereby enhancing line availability and overall efficiency.

In summary, the empirical results corroborate previous evidence: LIRL attains an average coverage of 92.78% with a mean performance loss below 7.5% across random noise and machine-failure scenarios, demonstrating strong OOD generalization and industrial-grade robustness.

Field study: real-world deployment in an elevator door-header factory

The LIRL scheduler was deployed in a commercial elevator door-header factory (Figure 6A) to (i) assign workpieces to two parallel production lines and (ii) orchestrate the three elemental robot operations in the hanging-plate assembly workcell. Each workpiece requires a fixed process plan with three sequential operations: O_1 (tension-spring-bracket assembly), O_2 (door-lock assembly), and O_3 (hanging-plate transfer). The policy was trained entirely offline using historical execution logs; no online exploration was conducted during deployment to ensure zero disruption to shop-floor operations. A detailed description of the hardware architecture is given in the Supplementary Materials.

In a pilot deployment, we evaluated an order comprising five workpieces. Under identical conditions of equipment, matched production environment, and personnel allocation, scheduling plans were generated and executed by both the existing rule-based scheduling system (serving as the baseline) and the proposed LIRL strategy. To support result re-reproducibility and facilitate extended statistical analysis, monitoring devices were deployed on the production line to systematically record the following data: (i) robot joint angles and power consumption under high-frequency sampling; (ii) the time required for the robot to complete each operation; and (iii) all constraint verification processes and associated interlock statuses. Primary outcomes are (i) makespan (s), (ii) aggregate electrical energy (kJ), and (iii) line utilization.

Schedules produced by the plant's rule-based dispatcher and by the proposed LIRL policy are executed and compared (Figure 6B and 6C, respectively). The LIRL schedule completes the order in 132.1 s, a 26.1% reduction relative to the 178.7 s baseline. Aggregate electrical energy drawn by the two robots drops from 312.1 kJ to 178.7 kJ, yielding a 42.7% saving. The rule-based logic keeps Line 0 saturated while leaving Line 1 idle for roughly one-third of the horizon (utilization 66.7%). LIRL raises utilization to 98.2% (Line 0) and 100% (Line 1), achieving a nearly balanced flow.

The above improvements originate from the multi-objective reward adopted during offline training, which jointly makespan, cumulative energy, and workload imbalance. The learned policy not only shortens production time but also smooths robot velocity profiles, thereby lowering energy demand without compromising throughput. Policy inference takes merely a few seconds, rendering the approach suitable for rolling-horizon rescheduling in discrete-part manufacturing.

The field trial demonstrates that an offline-trained LIRL scheduler can simultaneously improve makespan, energy efficiency, and line balance in a multi-line, multi-robot environment. These results substantiate the viability of reinforcement-learning-driven production control as a plug-and-play component for digital, low-carbon factories.

Simulation-based evaluations in other domains: Transportation and Micro-grids

The applicability of LIRL was examined in two representative CPS domains: transportation and micro-grids. Full scenario specifications are provided in the Supplementary Material.

Based on the quantitative results shown in Figure 7, we report the following findings. For clarity, the discussion is organized according to two canonical CPS tasks, i.e. urban traffic control and EV-charging micro-grids, followed by a cross-scenario synthesis.

Urban traffic control. Compared with the baselines, LIRL reduces the *network-wide average queue length* by 43.5% against PDQN and by 38.8% against HPPO. Meanwhile, the *system throughput* is boosted to 29.51×10^2 veh/h, i.e. $34.7 \times$ PDQN and $2.41 \times$ HPPO (Figure 7A.II). Throughout the entire evaluation horizon, LIRL introduces *zero* signal-phase conflicts and no green-time violations, whereas HPPO records 372.8 violations. PDQN is violation-free but with low efficiency. The results demonstrate that LIRL simultaneously mitigates congestion and maintains the required safety envelope, hence achieving an efficiency–safety dual optimum.

EV charging micro-grid

LIRL attains an average daily revenue of 12 456, outperforming PDQN and HPPO by 29.8% and 102%, respectively (Figure 7B.II). The *charger utilization ratio* and the *average output power* rise to 52.71% and 146.95 kW; both metrics exceed PDQN and HPPO, confirming LIRL's superior resource allocation under fluctuating loads. The *per-vehicle success rate* reaches




Figure 7. Results of the extended experiment. **A.I.** The urban road network with a 3×5 grid layout. **A.II** presents the performance comparison of the LIRL, PDQN, and HPPO algorithms using these metrics: Queue Length, average number of vehicles waiting at each intersection, sampled every 10 seconds; Throughput, average number of vehicles passing through each intersection, sampled every 10 seconds; Violations, total number of traffic rule violations across all intersections during the one-hour simulation. **B.I.** Overview of EV-charging station. **B.II** presents the performance comparison of the LIRL, PDQN, and HPPO algorithms using these metrics: Average Revenue, daily income, excluding costs from charger depreciation; Average Output Energy, total energy output per day; Average Charging Power, mean charging power over the day. Average Station Utilization, average utilization rate of the charging stations; Violation Rate, degree to which the scheduling decisions violate operational constraints. Success Rate, probability that an arriving vehicle successfully starts a charging session.

90.40% with zero violations of transformer capacity or current limits, in contrast to PDQN (24.77%) and HPPO (57.34%) violation rates. Hence, LIRL yields the best trade-off among profit, utilization, and operational safety.

Cross-scenario synthesis

By means of adaptive dual variables, LIRL balances optimization and feasibility, overcoming the classic efficiency–safety dilemma in RL. The zero-violation logs indicate that LIRL preserves constraint adherence during both learning and evaluation phases, a prerequisite for real-world deployment on critical infrastructures. Consistent superiority across heterogeneous tasks suggests that the LIRL framework can transfer to other energy and transport domains, positioning it as a promising candidate for large-scale CPS optimization.

Experimental evidence confirms that LIRL (i) maximizes system efficiency while *strictly* meeting operational constraints; (ii) generalizes across tasks of different scales and structures; and (iii) offers a compelling empirical basis for *safe-to-deploy reinforcement learning*.

Discussion

This work advances the state of cyber–physical cross-domain optimization by introducing logic-informed reinforcement learning, a generic framework that fuses declarative constraint reasoning with gradient-based policy learning. Through a projection operator that maps latent actions onto the dynamically constructed feasible manifold, LIRL converts standard

policy-gradient algorithm into a constraint-preserving optimizer for hybrid discrete-continuous action spaces. Extensive experiments on a representative robotic assembly line demonstrate that LIRL (i) improves makespan–energy trade-offs by 36.47–44.33% at most over hierarchical schedulers, (ii) surpasses recent constrained and hybrid-action RL baselines while eliminating all training-time violations, (iii) accelerates convergence by a factor of ten to forty, and (iv) maintains robustness under stochastic task durations and machine faults. The findings have several implications and limitations that merit discussion.

The projection mechanism decouples feasibility from reward shaping, thereby removing a major practical impediment to deploying RL in safety-critical CPS. Plant operators no longer need to hand-craft intricate penalty functions or retune them whenever the environment changes; instead, they simply encode domain knowledge as first-order logic templates. Because these templates are declarative, the same RL backbone can be ported across domains with minimal engineering effort. Moreover, LIRL accelerates learning by avoiding invalid exploration. This property reduces the amount of simulated or real-world interaction data required to reach competent performance, shortening the commissioning time of RL-enabled industrial controllers. In the context of robotics, the guaranteed feasibility also eliminates safety cages or hardware interlocks during policy refinement, lowering deployment costs.

Despite these strengths, several limitations warrant attention. Our benchmark test used 100 products and 5 work units—a scale that already challenges traditional optimization methods but remains moderate compared to industrial settings involving hundreds of workstations. Current physical constraints are limited to first-order linear forms; incorporating more complex constraints (e.g., multi-robot collision avoidance) could incur substantial computational cost, possibly requiring incremental or sampling-based methods for real-time decision-making. Although empirically and locally we show that the projection operator preserves policy-gradient fixed points, a general convergence proof for non-convex hybrid manifolds is still open. Extending projected stochastic optimization to hybrid action spaces is a promising direction. The framework also assumes stationary constraints, yet real-world CPS often face changing regulations or hardware updates. Future work should explore online logic synthesis and lifelong constraint learning for adaptive compliance.

Methods

This section details the problem formalization, the logic-informed reinforcement-learning framework, the projection operator, and the experimental protocol. All implementation scripts, configuration files and raw data are released at <https://github.com/wanguangxi/LIRL-CPS> to ensure full reproducibility. For more implementation details, please refer to the supplementary materials.

Problem formulation

We cast CPS cross-domain optimization as a constrained Markov decision process (CMDP)

$$\mathbf{M} \triangleq \langle \mathbf{S}, \mathbf{A}, \mathbf{T}, r, \mathbf{C}, \gamma, \Phi(s, a), \mathbb{F}(s), \mathbf{z}, \Pi_s \rangle,$$

where state $s \in \mathbf{S}$ captures the joint cyber–physical status of the plant, action $a \in \mathbf{A}$ is hybrid, transition kernel $\mathbf{T}$ encodes stochastic process dynamics, $r(s, a)$ is the primary scalar reward, $\gamma \in (0, 1)$ the discount factor, and $\mathbf{C} \triangleq \{c_k(s, a) \leq 0\}_{k=1}^K$ denotes K hard constraints. $\mathbf{A} = \mathbf{A}^d \times \mathbf{A}^c$ is the Cartesian product of a finite discrete set $\mathbf{A}^d$ and a continuous box $\mathbf{A}^c \subset \mathbb{R}^p$. $\Phi(s, a) := \bigwedge_{k=1}^K \phi_k(s, a)$ – conjunctive set of K state-dependent constraints. Each ϕ_k is assumed to be twice continuously differentiable in its continuous arguments and can be evaluated in $O(1)$ time for the discrete part. $\mathbb{F}(s) := \{a \in \mathbf{A} \mid \Phi(s, a) = \text{True}\}$ – state-conditioned feasible set. By engineering design $\mathbb{F}(s) \neq \emptyset$ for every s . $\mathbf{z} \in \mathbb{R}^L$ – latent “united” action produced by the policy network; $L \ll |\mathbf{A}|$. $\Pi_s : \mathbb{R}^L \rightarrow \mathbb{F}(s)$ – state-conditioned projection operator; $a = \Pi_s(z)$. $\pi_\theta(z|s)$ – latent-space stochastic policy with parameters θ ; the executed policy is $\pi_\theta^\Pi(a|s) = \int \mathbb{I}\{\Pi_s(z) = a\} \pi_\theta(z|s) dz$. Our objective is to maximize the discounted return $J(\pi_\theta^\Pi) = \mathbb{E}_\pi[\sum \gamma^t r_t]$ subject to $c_k(s_t, a_t) \leq 0 \forall k, t$.

Projection

For any state s and latent vector z , Π_s is computed in two independent steps:

Discrete action projection. Let $u = f_d(z) \in \mathbb{R}^{|\mathbf{A}^d|}$ be logits decoded from z . We solve

$$a^d = \arg \max_{a \in \mathbf{A}^d} \langle u, a \rangle \quad \text{s.t.} \quad \phi_{\text{cap}}(s, a) = \text{True}, \phi_{\text{prec}}(s, a) = \text{True}, \quad (1)$$

where $\phi_{\text{cap}}, \phi_{\text{prec}}$ represent the capacity constraint and the precedence constraint respectively. Because the constraints are linear and totally unimodular, the $\arg \max$ is equivalent to a linear-assignment problem solvable in $O(|\mathbf{A}^d|^3)$ by the Hungarian algorithm.

Continuous action projection. Let $v = f_c(z) \in \mathbb{R}^P$ be preliminary spline parameters. We solve a strictly convex quadratic program

$$a^c = \arg \min_{x \in \mathbb{R}^P} \|x - v\|^2 \quad \text{s.t.} \quad \phi_{\text{kin}}(s, a^d, x) \leq 0, \phi_{\text{col}}(s, a^d, x) \leq 0, \quad (2)$$

where $\phi_{\text{kin}}, \phi_{\text{col}}$ represent the kinematic constraints and collision constraints of the robot respectively. Slater's condition holds by design, hence the KKT system has a unique solution; we employ operator splitting quadratic program²⁸, OSQP with average complexity $O(P^2)$.

Finally $\Pi_s(z) := (a^d, a^c)$. The detailed algorithms can be found in the Supplementary Material.

Training

Each episode runs for at most 500 steps. When the replay buffer holds more than 500 transitions, the agent updates both actor and critic for 20 gradient steps per environment step. Models are checkpointed every 10 episodes for later evaluation. Unless otherwise specified, all reported metrics are averaged over 10 random seeds to ensure statistical significance. The LIRL policy using a single NVIDIA RTX4090 GPU. The training process takes 0.5 to 2.8 hours depending on problem size.

Constraint satisfaction guarantee

Lemma 1 (Feasibility) For every state s and latent $\mathbf{z}$, the action $a = \Pi_s(\mathbf{z})$ belongs to $\mathbb{F}(s)$.

Proof The discrete step enforces all purely cyber constraints; the continuous QP is solved over the intersection of physical constraints. By construction each ϕ_k is satisfied, hence $\Phi(s, a) = \text{True}$ and $a \in \mathbb{F}(s)$. ■

Regularity of Π

Assumption 1 (Regular constraints)

- (a) Each continuous constraint $\phi_k(s, a^c)$ is C^2 and its gradient with respect to a^c is L_k -Lipschitz.
- (b) For every s the Jacobian of active constraints has full row rank at the optimum of the QP.

Proposition 1 (Differentiability) Under Assumption 1, the mapping $a^c = \Pi_c(s, z)$ is C^1 almost everywhere in z ; its Jacobian is given by implicit differentiation of the KKT conditions. The overall Π_s is therefore piecewise-smooth.

Corollary 1 (Lipschitz continuity) There exists $L_\Pi > 0$ such that $\|\Pi_s(z_1) - \Pi_s(z_2)\| \leq L_\Pi \|z_1 - z_2\|$ for all z_1, z_2 .

Compatibility with policy-gradient methods

We consider the episodic performance objective

$$J(\theta) = \mathbb{E}_{s_0, z_0: T-1} \left[\sum_{t=0}^{T-1} r(s_t, \Pi_{s_t}(z_t)) \right]. \quad (3)$$

The latent-space likelihood is $\pi_\theta(z_t | s_t)$. Using the standard score-function estimator yields the gradient

$$\nabla_\theta J(\theta) = \mathbb{E} \left[\sum_t \nabla_\theta \log \pi_\theta(z_t | s_t) \cdot G_t \right], \quad (4)$$

where G_t denotes the cumulative return from t . Importantly, Π appears only inside the return, not inside the score term, so the estimator remains unbiased provided Π is deterministic and measurable.

Theorem 1 (Preservation of stationary points) Assume

- $r(s, a)$ is bounded and C^1 in its continuous arguments;
- Π_s satisfies Lemma 1 and Proposition 1;
- The latent policy π_θ is C^1 and non-degenerate on $\mathbb{R}^L$.

Then any θ^* that is a stationary point of the projected objective $J(\theta)$ is also a stationary point of the original CMDP objective restricted to feasible actions, i.e., $\nabla_\theta J(\theta^*) = 0 \Rightarrow \theta^*$ corresponds to a policy whose executed distribution $\pi_{\theta^*}^\Pi$ cannot be improved within the feasible set by an infinitesimal latent variation.

Proof Because Π_s is Lipschitz and piecewise-smooth, we can interchange differentiation and expectation (dominated convergence). The chain rule gives

$$\nabla_\theta r(s, \Pi_s(z)) = \partial_a r(s, a)|_{a=\Pi_s(z)} \cdot \partial_\theta \Pi_s(z) + 0, \quad (5)$$

but $\partial_\theta \Pi_s(z) = \partial_z \Pi_s(z) \cdot \partial_\theta z$, with $\partial_z \Pi_s$ bounded by Proposition 1. Consequently, the projected policy-gradient has the same zero set as the gradient of the constrained objective defined directly over $\mathbb{F}$. ■

Convergence

Let $\hat{g}_t$ be the Monte-Carlo estimate of $\nabla_{\theta} J(\theta)$ at iteration t . With a learning rate sequence η_t satisfying $\sum \eta_t = \infty$, $\sum \eta_t^2 < \infty$, the update $\theta_{t+1} = \theta_t + \eta_t \hat{g}_t$ constitutes a Robbins–Monro stochastic approximation.

Corollary 2 (Almost-sure convergence) Under Theorem 1 and the usual variance-boundedness condition on $\hat{g}_t$, the parameter sequence $\{\theta_t\}$ converges almost surely to the set of stationary points of $J(\theta)$.

Remark 1 If the continuous constraints are further convex and Π_c reduces to Euclidean projection, convergence rate matches that of projected stochastic gradient ($O(1/\sqrt{t})$). Non-convex collision constraints break global guarantees, yet empirical results confirm stable convergence in all tested instances.

References

1. Zhang, K. *et al.* Advancements in industrial cyber-physical systems: An overview and perspectives. *IEEE Transactions on Ind. Informatics* **19**, 716–729 (2023).
2. Ye Yuan, W. Z. W. P. X. L. H.-T. Z. H. D., Xiuchuan Tang & Goncalves, J. Data driven discovery of cyber physical systems. *Nat. Commun.* 4894, DOI: [10.1038/s41467-019-12490-1](https://doi.org/10.1038/s41467-019-12490-1) (2019).
3. Feng, S., Cucuzzella, M., Bouman, T., Steg, L. & Scherpen, J. M. A. An integrated human–cyber–physical framework for control of microgrids. *IEEE Transactions on Smart Grid* **14**, 3388–3400 (2023).
4. Li, H. & Jin, S. The integration of macroscopic traffic optimization control and microscopic traffic flow model for mixed traffic: A cyber-physical system perspective. *IEEE Transactions on Intell. Transp. Syst.* **25**, 20151–20161 (2024).
5. Wan, G. & Zeng, P. Codesign of architecture, control, and scheduling of modular cyber-physical production systems for design space exploration. *IEEE Transactions on Ind. Informatics* **18**, 2287–2296 (2022).
6. Wang, F. Vehicle 5.0: From driverless vehicles for its in cps to metadriving for smart mobility in cpss. *IEEE Trans. Intell. Veh.* **8**, 3523–3526 (2023).
7. Cheng, T., Duan, T. & Dinavahi, V. R. Ecs-grid: Data-oriented real-time simulation platform for cyber-physical power systems. *IEEE Transactions on Ind. Informatics* **19**, 11128–11138 (2023).
8. Hu, C.-L., Wang, L., Chen, M.-L. & Pei, C. A real-time interactive decision-making and control framework for complex cyber-physical-human systems. *Annu. Rev. Control.* **57**, 100938 (2024).
9. Babayiğit, B. & Abubaker, M. Industrial internet of things: A review of improvements over traditional scada systems for industrial automation. *IEEE Syst. J.* **18**, 120–133 (2024).
10. Huang, D. *et al.* Consensus-based distributed quantum decomposition algorithm for security-constrained unit commitment considering optimal transmission switching. *Adv. Quantum Technol.* e2500241, DOI: <https://doi.org/10.1002/qute.202500241>.
11. Xiong, J. *et al.* Parametrized deep q-networks learning: Reinforcement learning with discrete-continuous hybrid action space. *ArXiv* **abs/1810.06394** (2018).
12. Fan, Z., Su, R., Zhang, W. & Yu, Y. Hybrid actor-critic reinforcement learning in parameterized action space. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence, IJCAI'19*, 2279–2285 (AAAI Press, 2019).
13. Mnih, V. *et al.* Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015).
14. DeepSeek-AI *et al.* Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *Nature* **645**, 633–638 (2025).
15. Schrittwieser, J. *et al.* Mastering atari, go, chess and shogi by planning with a learned model. *Nature* **588**, 604 – 609 (2019).
16. Song, W., Chen, X., Li, Q. & Cao, Z. Flexible job-shop scheduling via graph neural network and deep reinforcement learning. *IEEE Transactions on Ind. Informatics* **19**, 1600–1610, DOI: [10.1109/TII.2022.3189725](https://doi.org/10.1109/TII.2022.3189725) (2023).
17. Degraeve, J. *et al.* Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature* **602**, 414 – 419 (2022).
18. Triantafyllidis, E., Acero, F., Liu, Z. & Li, Z. Hybrid hierarchical learning for solving complex sequential tasks using the robotic manipulation network roman. *Nat. Mach. Intell.* **5**, 991 – 1005 (2023).
19. Wang, Y., Wu, J., He, H., Wei, Z. & Sun, F. Data-driven energy management for electric vehicles using offline reinforcement learning. *Nat. Commun.* **16** (2025).

20. Singh, R., Gupta, A. & Shroff, N. B. Learning in constrained markov decision processes. *IEEE Transactions on Control. Netw. Syst.* **10**, 441–453, DOI: [10.1109/TCNS.2022.3203361](https://doi.org/10.1109/TCNS.2022.3203361) (2023).
21. Ming, J., Bargmann, D., Cao, H. & Caccamo, M. Flexible gear assembly with visual servoing and force feedback. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 8276–8282, DOI: [10.1109/IROS55552.2023.10341833](https://doi.org/10.1109/IROS55552.2023.10341833) (2023).
22. Li, G. *et al.* Development of a manufacturing system for gear assembly using collaborative robots. In *2021 IEEE International Conference on Mechatronics and Automation (ICMA)*, 22–27, DOI: [10.1109/ICMA52036.2021.9512631](https://doi.org/10.1109/ICMA52036.2021.9512631) (2021).
23. Li, B. *et al.* Hyar: Addressing discrete-continuous action reinforcement learning via hybrid action representation. *10th Int. Conf. on Learn. Represent. ICLR 2022 Virtual, Online* (2022).
24. Achiam, J. & Amodei, D. Benchmarking safe exploration in deep reinforcement learning (2019).
25. Achiam, J., Held, D., Tamar, A. & Abbeel, P. Constrained policy optimization. *ArXiv* **abs/1705.10528** (2017).
26. Huang, S. & Ontan'on, S. A closer look at invalid action masking in policy gradient algorithms. *ArXiv* **abs/2006.14171** (2020).
27. Lillicrap, T. P. *et al.* Continuous control with deep reinforcement learning. *CoRR* **abs/1509.02971** (2015).
28. Stellato, B., Banjac, G., Goulart, P., Bemporad, A. & Boyd, S. Osqp: An operator splitting solver for quadratic programs. In *2018 UKACC 12th International Conference on Control (CONTROL)*, 339–339, DOI: [10.1109/CONTROL.2018.8516834](https://doi.org/10.1109/CONTROL.2018.8516834) (2018).

Data and code availability

All source code of LIRL, logic files and results are on <https://github.com/wanguangxi/LIRL-CPS>.

Acknowledgments

This work was supported by the National Key Research, and Development Program of China [2024YFB4711103], National Natural Science Foundation of China [U24A20277, 92267205, 62503466, 92467301, 92367301], Natural Science Foundation of Liaoning Province [2024-MSBA-83, 2025-MS-085], Fundamental Research Project of SIA [2024JC1K10], and the State Key Laboratory of Robotics and Intelligent Systems of China [2025-Z12].

Author contributions statement

Guangxi Wan and Peng Zeng conceived the experiments, Yiyang Liu, Xiaoting Dong and Hongfei Bai conducted the experiments, Shijie Cui, Qingwei Dong Dong Li, and Chunhe Song analyzed the results. All authors reviewed the manuscript.