

SHaRe-RL: Structured, Interactive Reinforcement Learning for Contact-Rich Industrial Assembly Tasks

Jannick Stranghöner¹, Philipp Hartmann¹, Lisa-Marie Weigelt¹,
 Marco Braun², Sebastian Wrede^{1,2}, Klaus Neumann^{1,2}

Abstract—High-mix low-volume (HMLV) industrial assembly, common in small and medium-sized enterprises (SMEs), requires the same precision, safety, and reliability as high-volume automation while remaining flexible to product variation and environmental uncertainty. Current robotic systems struggle to meet these demands. Manual programming is brittle and costly to adapt, while learning-based methods suffer from poor sample efficiency and unsafe exploration in contact-rich tasks. To address this, we present SHaRe-RL, a reinforcement learning framework that leverages multiple sources of prior knowledge. By (i) structuring skills into manipulation primitives, (ii) incorporating human demonstrations and online corrections, and (iii) bounding interaction forces with per-axis compliance, SHaRe-RL enables efficient and safe online learning for long-horizon, contact-rich industrial assembly tasks. Experiments on the insertion of industrial Harting connector modules with 0.2–0.4 mm clearance show reliable learning within practical wall-clock budget and improved performance over an unstructured human-in-the-loop RL baseline. We further show that the learned policy generalizes to previously unseen connector variants. Overall, our results show that process expertise alone can effectively guide real-world RL, making deployment safer, more robust, and economically viable for industrial assembly. Source code and demonstration videos are available at <https://share-rl.github.io/share-rl.io/>

I. INTRODUCTION

High-mix low-volume (HMLV) industrial assembly, which is ubiquitous in small and medium-sized enterprises (SMEs), remains one of the most difficult challenges for autonomous robots [1]. These settings demand near-zero failure rates, high precision across diverse product variants, and strict compliance with safety regulations, all while operating under inherent environmental uncertainty [2], [3]. The challenge is exacerbated by the fact that SMEs typically lack automation experts, yet depend on production flexibility to remain competitive. Consequently, such tasks are still carried out almost exclusively by human operators. Growing labor shortages further intensify the need for robots that can adapt to changing assembly requirements, but today’s industrial systems remain notoriously inflexible.

Reinforcement learning (RL) promises adaptable robots that can handle uncertainty and even outperform hand-designed controllers or teleoperation [4]. However, deploying RL in real-world industrial settings remains difficult due to sample inefficiency, under- or overspecified reward functions,

¹This work was supported by the Fraunhofer Internal Programs under Grant No. SME 40-09551.

¹CITEC, Faculty of Technology, Bielefeld University, Germany.

²Fraunhofer IOSB-INA, Lemgo, Germany.

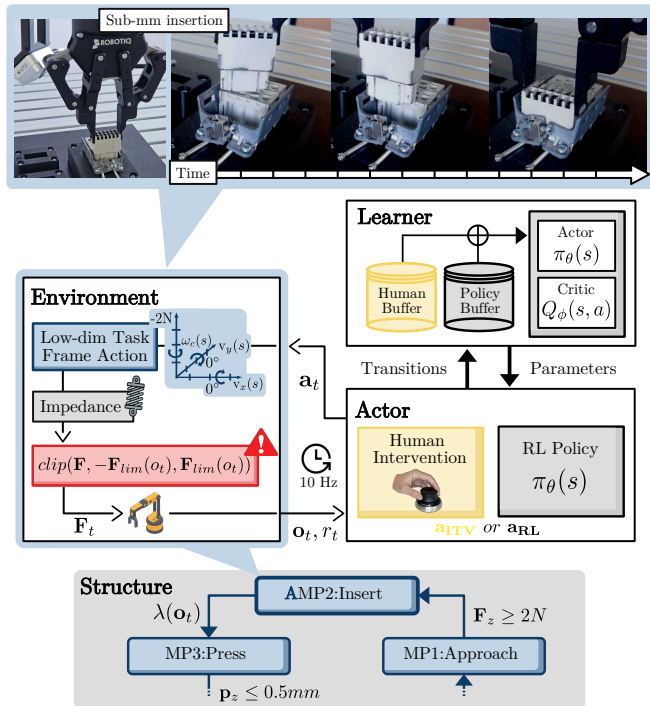


Fig. 1: **Structure, interaction and safety.** SHaRe-RL is a hybrid approach that combines a hand-designed primitive sequence with interactive learning in selected primitives and adaptive per-axis force limits for bounded contact forces.

and the risks of unsafe exploration in contact-rich interactions [5], [6].

In order to close this gap, prior work has emphasized the importance of injecting domain knowledge into RL [7]. Notably, two powerful forms of prior knowledge commonly exist in industrial settings: process structure, captured by production experts during engineering time through established formalisms [8], and operator expertise, provided through demonstrations and online corrections [9]. Each of these has proven effective in isolation [3], [10], but they have not been combined with modern, sample-efficient RL under real-world safety constraints.

We present SHaRe-RL, a Structured, Human-in-the-loop, and Real-world RL framework that integrates these sources of prior knowledge with a simple but effective compliance mechanism. Fig. 1 illustrates how SHaRe-RL embeds sample-efficient, interactive RL (yellow) within established task abstractions such as manipulation primitives (blue), while ensuring safety through adaptive force limits (red).

SHaRe-RL scales to complex, vision-based assembly by guiding exploration to meaningful subspaces and focusing learning where it is most beneficial. Through demonstrations and online corrections, SHaRe-RL exploits operator expertise to accelerate exploration, while RL itself is able to improve beyond suboptimal human behavior. This hybrid approach preserves interpretability, supports effective human-robot collaboration, and enables faster, safer, and more economically viable training, moving RL a step closer to broad industrial acceptance. Our contributions are as follows:

- We introduce SHaRe-RL as the first system to combine manipulation primitive nets with human-in-the-loop, off-policy reinforcement learning.
- We formulate an adaptive per-axis compliance mechanism that provably bounds interaction forces during exploration while preserving free-space dynamics.
- We demonstrate empirically that SHaRe-RL achieves robust performance on a sparse-reward industrial insertion task with sub-millimeter tolerances within three hours of real-world interaction, beating the state-of-the-art and even surpassing the cycle time of a skilled human demonstrator.
- We show generalization to previously unseen connector variants and thoroughly validate SHaRe-RL through ablations and analysis.
- We implement SHaRe-RL as an extension of the open-source LeRobot framework [11], making it compatible with a wide range of low-cost manipulators.

II. RELATED WORK

Contact-rich assembly is difficult to solve end-to-end with reinforcement learning, as large state–action spaces, sample efficiency, and strict safety constraints hinder direct trial-and-error learning [2]. Robotics has therefore long relied on abstractions to manage complexity, such as manipulation primitives (MPs) and the task frame formalism [12]–[14]. These abstractions can be composed into higher-level structures called Manipulation Primitive Nets (MP-Nets) [8]. Early formulations that combine abstractions with learning-based methods treat primitives as discrete choices with small parameter sets to be optimized [15], [16]. Their process-level coordination can also be learned via hierarchical RL [17]. More recent approaches explore neural-network based primitives, first in isolation for compliant skills such as cutting [18], and later within assembly sequences [3]. While these approaches demonstrate that task structure can drastically reduce exploration complexity, they typically assume on-policy updates, low-dimensional state features and carefully engineered reward functions. This limits their applicability to realistic industrial tasks with high-dimensional input and sparse rewards.

In parallel, advances in sample-efficient RL have shown that offline demonstrations and online interventions can drastically accelerate learning [4]–[6]. Off-policy RL with demonstrations [4], [19], [20], model-based approaches [21], [22], interactive or residual imitation learning [23]–[25], and

distributed frameworks like SERL [26] and its human-in-the-loop extension HIL-SERL [10] demonstrate that human input can make sparse-reward, real-world tasks tractable. However, these approaches are typically trained and evaluated as isolated skills rather than as components of larger sequences and rarely exploit the forms of prior knowledge that are widely available in practice. Our contribution is to embed the same class of sample-efficient, human-in-the-loop RL into structured assembly processes.

Traditional control approaches often achieve robustness but remain conservative and depend heavily on engineering expertise [27], whereas learning-based controllers promise greater adaptability but face challenges in safety and interpretability. Here, classical optimization-based safety layers, such as control barrier functions (CBFs), quadratic program (QP) filters, and model predictive control (MPC) shields, can enforce constraints during execution, but trade off model accuracy against computational overhead [28]–[30]. Recent work instead learns variable-compliance schemes that adapt robot stiffness through RL, but lower stiffness alone cannot prevent unbounded contact forces when the policy issues large set-points, requiring careful tuning to avoid [31]–[33]. In contrast, deterministic schemes that directly expose physically meaningful limits offer interpretability and ease of deployment. Our method contributes to this line by introducing per-axis adaptive force limits, which provably bound contact forces while preserving free-space dynamics, enabling RL to train safely in contact-rich settings.

III. METHOD

SHaRe-RL combines three complementary sources of domain knowledge into one asynchronous learning pipeline. In this section we first show how human demonstrations and real-time operator interventions are incorporated into an off-policy RL backbone; we then explain how task–frame manipulation primitives focus exploration on the phases that matter most; and we conclude with a simple yet provably safe compliance layer that bounds interaction forces.

A. Operator Guidance: Human-in-the-Loop RL

Robotic reinforcement learning tasks can be formulated as a finite-horizon Markov Decision Process (MDP) where the objective is to learn a policy which maximizes the expected sum of discounted future rewards [34]. We adopt a distributed human-in-the-loop RL training loop similar to HIL-SERL [10], which has proven effective for real-world RL with sparse rewards.

Before training begins, a skilled operator collects a set of teleoperated trajectories with a six–DoF SpaceMouse, which are stored in a demonstration buffer. During training, the operator can press a button to take control of the robot at any time. While the button is held, the policy action is replaced by the human SpaceMouse action. This biases data collection toward successful and safe behaviors and helps the policy escape poor local optima [10]. Beyond corrections, the operator can also deliberately induce rare but relevant failure cases, forcing the policy to learn robust recovery

behavior. In practice, operators intervene frequently during the early stages of training and reduce their involvement as performance improves, with faster convergence when corrections are short and specific while otherwise letting the policy explore on its own.

At the system level, an actor process executes the current policy on the robot, generates experience, and streams it to the learner. The learner process samples mini-batches equally from the demonstration buffer and an on-policy buffer to perform gradient updates on the RLPD loss [20]. RLPD is a regularized variant of SAC [35] designed for sample efficiency and integration of prior data without offline pretraining. Updated network weights are streamed back to the actor at fixed time intervals.

B. Task-Level Priors: TFF and MP-Nets

Human-in-the-loop corrections alone are not sufficient to enable long-horizon assembly, as the search space quickly becomes intractable without additional structure. A widely used strategy to add such structure is to decompose complex processes into manipulation primitives (MPs), each representing a fundamental capability of the manipulator [36]. Finkemeyer et al. [8] formalize MPs as the following tuple

$$\text{MP} := \{\text{HM}, \tau, \lambda\},$$

where HM describes a hybrid force- or position control policy for each translational and rotational direction relative to a task frame, τ is a tool command and λ is a boolean predicate on sensor values that defines a stop condition [12], [13]. Importantly, the type of controller can differ in all directions. Manipulation Primitive Nets (MP-Nets) extend this abstraction by composing MPs into state machines, where transitions are triggered by stop conditions [8].

Classical MPs are executed with fixed control setpoints. Extending this notion, adaptive MPs (AMPs) replace selected setpoints by policy-controlled variables [3]. If n axes are controlled by the policy, the corresponding action space is n -dimensional. The operator’s SpaceMouse is likewise constrained to these n axes, ensuring that both demonstrations and interventions remain aligned with the degrees of freedom available to the policy. This formulation focuses exploration on task-relevant directions and high-reward regions of the state-space. For example, in peg-in-hole insertion, the coarse approach is executed with fixed setpoints, while learning is reserved to handle uncertainty during fine alignment.

SHaRe-RL uses end-to-end demonstrations as a complementary source of prior knowledge to improve the sample efficiency of AMPs. Instead of training individual MPs in isolation with their own reset strategy, both demonstrations and online rollouts traverse the entire MP-Net from initial to terminal state. During data collection the operator teleoperates adaptive primitives with the SpaceMouse, while non-adaptive primitives run with fixed set-points; all predefined stop conditions are active. When a stop event is difficult to encode analytically, the operator can signal the end of a primitive with a button press. These labels are used to train

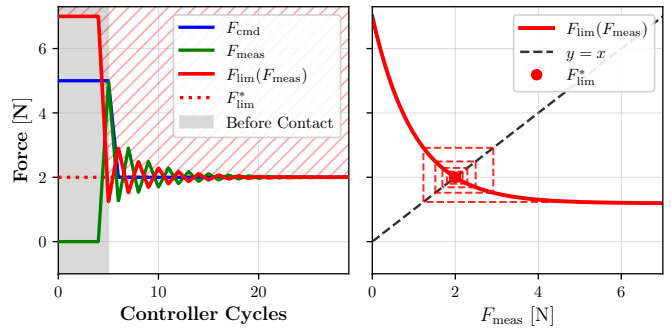


Fig. 2: Adaptive force limits under ideal conditions Left: Time response of a stable adaptive limit. Right: Phase-plane cobweb plot illustrating the recurrence in Eq.2.

a lightweight binary classifier that later triggers the same transition autonomously, which is common practice in real-world RL [37].

This design has three advantages. First, MP-Nets can explicitly encode reset strategies, such that the same graph that executes a task also returns the system to a well-defined start state. In our experiments, the MP-Net includes an explicit reset primitive from the terminal back to the start state, enabling fully automatic, reset-free rollouts without manual intervention [38]. Second, our software stack supports jointly optimizing multiple AMPs within one MP-Net, though we focus on a single AMP in this work and leave multi-primitive joint optimization to future work. Third, the initial state distribution of each AMP is correctly induced by its predecessors. Specifically, uncertainty introduced by upstream actions, such as variability in a grasp, must be handled robustly by downstream primitives. Training with “atomic” reset strategies would alter these distributions and yield policies that generalize poorly at execution time. Training with full MP-Net rollouts instead exposes policies to the exact state distributions they will encounter, improving robustness during deployment.

C. Adaptive Safety Limits

Exploration in contact-rich assembly can easily produce large, destructive contact forces. Beyond the risk of hardware damage, such events degrade sample efficiency by aborting episodes early and biasing trajectory distributions. To address this, we introduce adaptive per-axis force limits as a deterministic safety layer for compliant control. This mechanism enforces physically meaningful force bounds under contact without compromising responsiveness in free-space.

Let $F_{\text{cmd},t}$ denote the raw output of the high-level controller or policy, and $F_{\text{meas},t}$ the measured contact force at time step t . Given a static hardware bound $F_{\text{max}} > 0$, we define a time-varying adaptive $F_{\text{lim},t}$

$$F_{\text{lim},t}(F_{\text{meas},t}) = \alpha(|F_{\text{meas},t}|) F_{\text{max}} > F_{\text{cmd},t}, \quad (1)$$

where $\alpha : \mathbb{R}_{\geq 0} \rightarrow (0, 1]$ is a monotone decreasing map with $\alpha(0) = 1$. This limit is large during free-space motion, and contracts once contact is detected ($F_{\text{meas},t} > 0$).




Fig. 3: **Experimental setup and task.** Top: UR3e robot and Harting HanDD industrial connector used for evaluation. Bottom: Cropped policy input over time.

Assuming rigid-body contact ($F_{\text{meas},t+1} = F_{\text{cmd},t}$), the closed loop evolves as

$$F_{\text{meas},t+1} = F_{\text{lim},t} = \alpha(|F_{\text{meas},t}|) F_{\text{max}} \quad (2)$$

The long-term behavior of this recurrence is governed by its fixed point F_{lim}^* , which corresponds to the steady-state force established under contact. Convergence to F_{lim}^* is guaranteed whenever the recurrence is locally stable, i.e., it satisfies

$$|\alpha'(F_{\text{lim}}^*)| F_{\text{max}} < 1 \quad (3)$$

A practical choice for α is an exponential decay with minimal scaling α_{min} :

$$\alpha(F) = \alpha_{\text{min}} + (1 - \alpha_{\text{min}}) \exp(-F/\theta). \quad (4)$$

This form is smooth, monotone, and admits a unique fixed point, which can be set deterministically by solving for the corresponding decay parameter θ^* given F_{max} and a desired F_{lim}^* . Stability then reduces to a simple inequality on α_{min} , which is conceptually similar to contraction-based verification of dynamical systems [39]. Fig. 2 illustrates the idealized time and phase-plane response of a stable adaptive force limit. Large commands are allowed in free space, but once F_{meas} exceeds zero the limit shrinks and both commanded and measured forces settle at F_{lim}^* without manual tuning.

IV. CONNECTOR ASSEMBLY TASK

We evaluate our approach on the insertion of industrial Harting Han-Modular connectors, a representative contact-rich assembly task with tolerances of only 0.2–0.4 mm. The task resembles peg-in-hole insertion but involves non-cylindrical geometry and a final snap-in stage. While the learning problem itself is concentrated in the fine insertion phase, reliable real-world training benefits from structuring

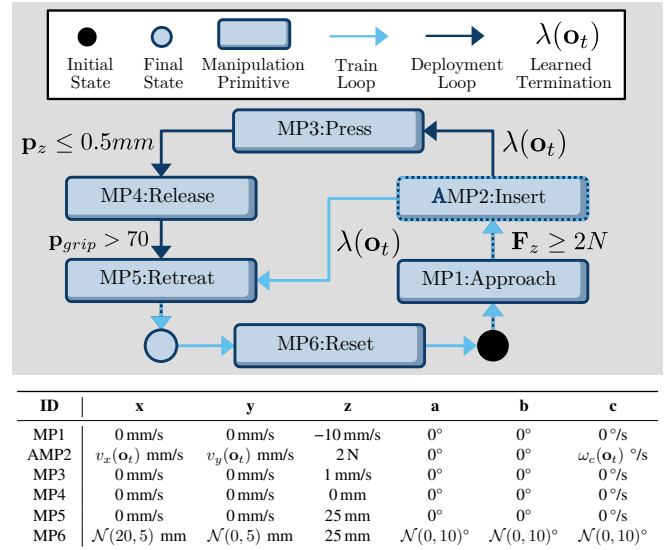


Fig. 4: **MP-Net for the HanDD connector insertion.** Light-blue and dark-blue arrows indicate the train and deployment loop respectively. The learned termination condition $\lambda(\mathbf{o}_t)$ and the randomized reset MP6 close the loop for repeated trials. The table lists the per-axis controller setpoints for each primitive.

the overall process into a small number of manipulation primitives. Scripted primitives handle the repeatable parts of the process (approach, coarse alignment and post-contact actions), while learning is restricted to the uncertain insertion where it is most beneficial. Fig. 3 illustrates the setup, and Fig. 4 summarizes the MP-Net used in our experiments, including an explicit reset primitive (MP6) that returns the robot from terminal back to the start state and enables fully automatic rollouts.

A. Robot and Controller

Experiments are conducted on the UR3e depicted in Fig. 3, which is equipped with an integrated six-axis force/torque sensor at the wrist, an ego-centric Intel RealSense D405 RGB-D camera, and custom 3D-printed gripper fingers. The control hierarchy consists of three layers. At the top, desired force or velocity setpoints in task frame axes are updated at 10 Hz. These targets are processed by a mid-level Cartesian impedance loop running at 500 Hz, which translates them into wrench commands for the robot’s proprietary force controller. This formulation enables us to apply per-axis force limits directly to the commanded wrench before it reaches the hardware. In free space, the maximum limits are set to 30 N for translation and 10 N m for rotation, while under contact the adaptive scaling settles at 5 N and 0.5 N m, respectively, ensuring safety during exploration.

B. MDP Formulation

The connector insertion is modeled as a finite-horizon MDP. To provide a clearer formalization, we decompose this into the observation and action space, reset mechanism, and reward function.

Observation Space: The observation at time t is defined as

$$\mathbf{o}_t = (I_t, \mathbf{x}_{t-1:t}, \mathbf{F}_{t-1:t}), \quad (5)$$

where $I_t \in \mathbb{R}^{128 \times 128 \times 3}$ is the wrist-mounted RGB image, $\mathbf{x}_{t-1:t} \in \mathbb{R}^{2 \times 12}$ are stacked end-effector poses and velocities, and $\mathbf{F}_{t-1:t} \in \mathbb{R}^{2 \times 6}$ are stacked wrenches. Frame stacking compensates for partial observability and non-Markovian effects introduced by the compliant controller. To prevent overfitting to a fixed geometry, absolute positions along learned axes are excluded from the state input. Images are cropped and downsampled to a local view, augmented with random shifts, and encoded with a frozen, pretrained ResNet10 [40], [41]. The resulting embeddings are concatenated with proprioceptive and wrench features before being processed by the policy and critic networks. Both networks are implemented as two-layer multilayer perceptrons (MLPs) with 256 hidden units per layer and SiLU activations [42]. The policy outputs the parameters of a diagonal Gaussian, which is squashed with \tanh to enforce action limits, following the standard practice from Soft Actor-Critic [35]. Unless stated otherwise, all methods in Table I receive the observation \mathbf{o}_t and use the same network architecture, preprocessing, and augmentation pipeline.

Action Space: The action space depends heavily on the structural priors provided to the agent. For our structured variants, the MDP is defined solely over the insertion primitive. Consequently, the action space is restricted to Cartesian velocity setpoints for the x -, y -, and c -axes. In contrast, the unstructured baselines must solve the insertion end-to-end without task-level priors, requiring them to act in the full six-DoF Cartesian space.

Reset: All methods share the same scripted reset mechanism and episode horizon. Compared to prior work [3], we expose an additional rotational degree of freedom around the connector’s c -axis and increase initial pose uncertainty. At every episode start, Gaussian noise is applied independently to each axis: ($\sigma_{x,y} = 5$ mm) in translation and ($\sigma = 10^\circ$) in each rotation axis. To simulate systematic calibration errors, the initial distribution is offset by 2 cm along the x -axis. As a result, the policy must traverse a longer distance before fine alignment, increasing the maximum episode length from 3 s to 9 s (90 steps).

Reward Function: Finally, the learning objective is defined by either a dense shaping reward based on vertical distance to the target, or a sparse binary success classifier trained offline:

$$r_t^{\text{dense}} = -\|z_t - z_{\text{goal}}\|_2, \quad r_t^{\text{sparse}} = \lambda(\mathbf{o}_t),$$

The classifier $\lambda(\mathbf{o}_t)$ uses the same network architecture as the policy and is trained offline before RL on a mix of 20 positive and 20 negative demonstrations. Since many industrial success signals are inherently sparse, and sparse classifiers tend to generalize better under uncertainty [43], we evaluate on r_t^{sparse} for our main experiments and use r_t^{dense} for ablations.

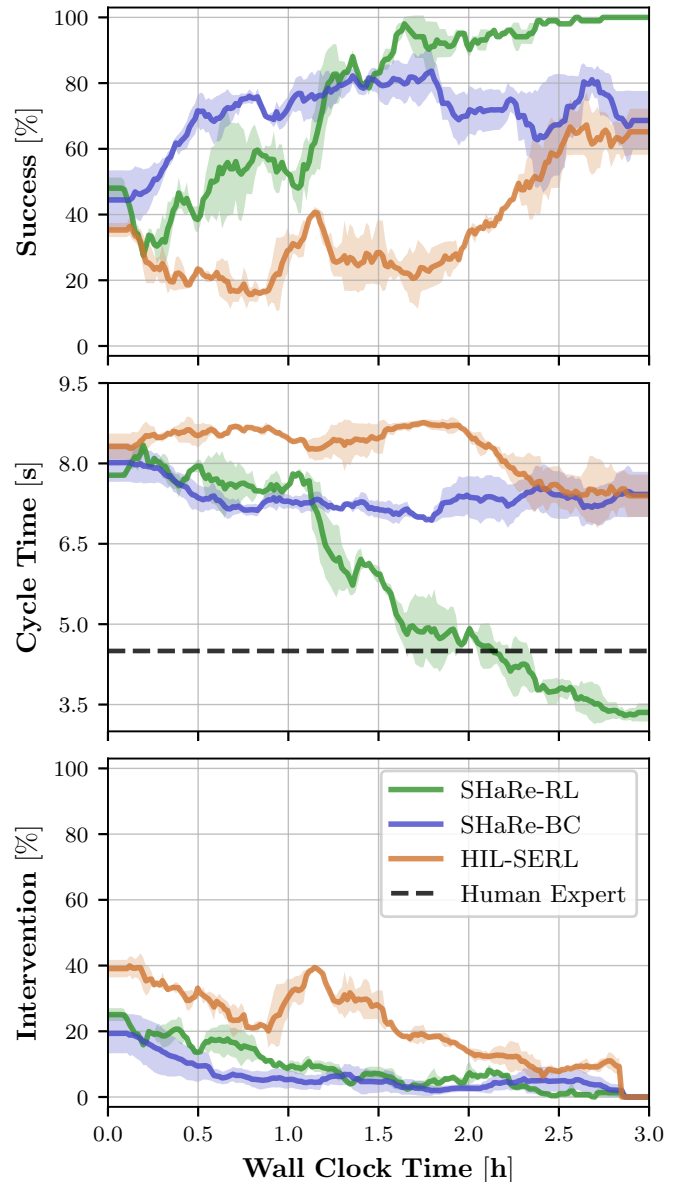


Fig. 5: **Main comparison of learning approaches.** Success rate, cycle time, and intervention rate vs wall-clock time (mean \pm s.e.m., 2 seeds).

V. EXPERIMENTAL RESULTS

A key premise of our method is that task-level motion primitives drastically reduce the size of the effective state-action space, and thereby lower the overall complexity of the learning problem. We evaluate this hypothesis by comparing structured interactive RL (SHaRe-RL) to (i) an unstructured human-in-the-loop RL baseline (HIL-SERL) that does not use MPs, and (ii) a structured interactive behavior cloning baseline (SHaRe-BC) similar to HG-Dagger [25]. In addition, we include pure RL baselines and structured ablations to isolate the role of demonstrations, interventions, and vision. We compare all methods over two training runs using a fixed wall-clock budget of three hours on a single NVIDIA RTX 3060 Ti workstation. Each variant is initialized with

| Method | Success [%] | Time [s] | Structure | Demos | Interventions |
|--------------------------------------|-------------|------------|-----------|-------|---------------|
| SAC _{dense} | 5 | 8.7 | ✗ | ✗ | ✗ |
| HIL-SERL | 65 | 7.4 | ✗ | ✓ | ✓ |
| SHaRe-BC | 80 | 7.4 | ✓ | ✓ | ✓ |
| SHaRe-RL | 100 | 3.5 | ✓ | ✓ | ✓ |
| SHaRe-RL _{PPO} | 5 | 8.8 | ✓ | ✗ | ✗ |
| SHaRe-RL _{no-demos} | 35 | 7.9 | ✓ | ✗ | ✗ |
| SHaRe-RL _{no-interventions} | 45 | 7.5 | ✓ | ✓ | ✗ |
| SHaRe-RL _{no-vision} | 10 | 8.8 | ✓ | ✓ | ✓ |
| Human Expert | 100 | 4.5 | - | - | - |
| Random Policy | 5 | 8.9 | - | - | - |

TABLE I: Main comparison across methods. SHaRe-RL variants use the same off-policy RL backbone but differ in which sources of prior knowledge are used. HIL-SERL represents human-in-the-loop RL without task-level priors, while SHaRe-BC represents structured interactive imitation without reinforcement learning.

20 teleoperated demonstrations and evaluated at the end of training using 20 autonomous rollouts without interventions. During training, we track success rate, cycle time, and intervention rate using a rolling average. The policy executes at 10 Hz; cycle time is measured from the beginning of a rollout after reset until the success condition is met or the episode times out. For structured methods, this definition includes the approach primitive and is therefore comparable across structured and unstructured methods.

A. Main Results

Fig. 5 and Table I summarize the main comparison. SHaRe-RL achieves the best overall performance, reaching 100% success with a 3.5 s cycle time, and thus outperforming both the human expert (4.5 s at 100%) and all learning baselines. Importantly, SHaRe-RL also improves substantially over the unstructured human-in-the-loop RL baseline HIL-SERL, which reaches 65% success at 7.4 s. This gap directly quantifies the practical benefit of adding structure to contact-rich real-world RL. SHaRe-BC quickly reaches strong performance (80% at 7.4 s), but plateaus near the quality of its demonstrations, whereas SHaRe-RL continues to improve throughout the full training budget. This reflects the core strengths of RL over BC: human demonstrations can be suboptimal due to pauses, conservative behavior or operator tendencies. For example, we observed that humans typically align each axis sequentially when using a space-mouse, resulting in longer trajectories. BC performance is bounded by the quality of the demonstrations, while RL can use dynamic programming to learn from both successes and failures. This allows SHaRe-RL to self-correct, discover behaviors that are hard to demonstrate, such as simultaneous multi-axis alignment with optimal dynamics, and ultimately exceed human performance. This is consistent with theoretical results [44].

We also evaluate intervention rates during training to measure the shift toward autonomous execution. Specifically, we compute the ratio of intervened timesteps to total timesteps per episode, averaged over a 20-episode window. As shown in Fig. 5, SHaRe-RL’s reliance on human corrections steadily

decreases and reaches zero by the end of training. Qualitatively, we observe a shift from long, sparse interventions early in training to brief, targeted corrections as the policy matures.

B. Ablation Study

The ablations in Table I disentangle the contributions of demonstrations, interventions, and vision. Removing online corrections (SHaRe-RL_{no-interventions}) reduces performance to 45% success, while further removing offline demonstrations (SHaRe-RL_{no-demos}) only achieves 35%, confirming that both sources of human guidance are important within the three-hour budget. The on-policy variant (SHaRe-RL_{PPO}) [45] reaches just 5% success after the full training budget, highlighting its inability to handle sparse rewards and vision input. Removing vision (SHaRe-RL_{no-vision}) collapses performance to 10%, indicating that image feedback is critical for resolving geometric ambiguity, in particular around the connector’s *c*-axis, which cannot be inferred reliably from force and velocity alone under the pose uncertainties considered.

To better understand the role of task and operator priors, we additionally visualize state-space occupancy during training in Fig. 7. We compare strictly autonomous rollouts of SHaRe-RL to an unstructured dense-reward baseline (SAC_{dense}) that is trained end-to-end without primitives, demonstrations, or interventions. Although the dense reward provides a continuous learning signal along the insertion direction, uncertainty in orientation and contact geometry renders this signal largely uninformative for alignment. As a result, SAC_{dense} explores a much larger region but remains centered above the opening and rarely enters the narrow funnel required for insertion. The policy fails to align the connector reliably, as the reward landscape is largely flat along rotational axes. In contrast, SHaRe-RL concentrates occupancy in the successful region even under sparse rewards, indicating that well-chosen priors can be

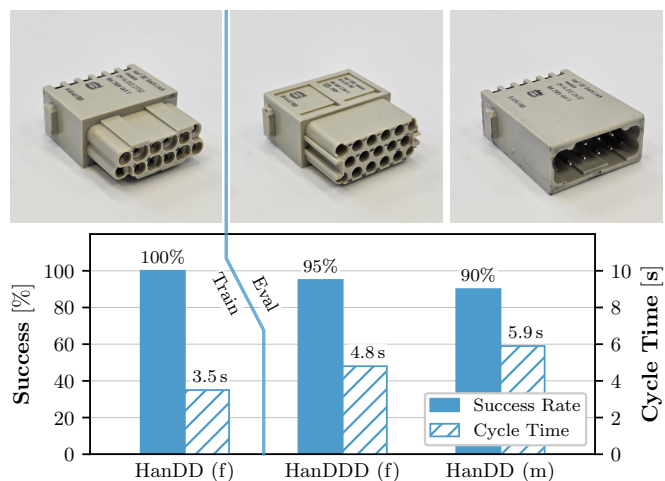


Fig. 6: Generalization across novel connectors. A policy trained on HanDD (f) is evaluated without finetuning on two previously unseen connector variants. We report success rate and cycle time over 20 autonomous rollouts per connector.

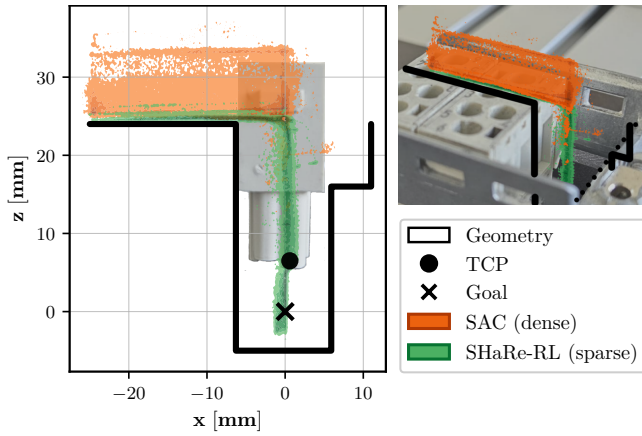


Fig. 7: **State-space occupancy during training.** Fully autonomous rollouts projected onto a 2D x - z slice of the workspace. SAC (dense), trained end-to-end without task structure, explores broadly but rarely enters the narrow funnel that leads to successful insertion. SHaRe-RL (sparse) concentrates occupancy in goal-relevant regions.

more effective than hand-crafted shaping rewards for guiding exploration in contact-rich settings.

C. Zero-Shot Generalization

We further evaluate the zero-shot transfer capabilities of the learned policy by deploying it, without any finetuning, on two previously unseen connector variants (Fig. 6). Although the average cycle time increases as the geometry and contact dynamics change, the success rate remains robust. This demonstrates that SHaRe-RL learns a generalized insertion strategy rather than memorizing a specific spatial configuration. Prior work has shown that transferring policies across connector geometries is difficult when relying solely on force feedback [3], supporting the necessity of image-based observations for robust assembly.

D. Adaptive Force Limits

For a task-level comparison, we compare our adaptive limits to the HIL-SERL-style limits that bound the impedance reference error and thus bound the maximum wrench at all times [10]. During demonstration collection, adaptive limits achieve 100% success with 4.5 s average cycle time, compared to 90% success and 5.2 s with HIL-SERL limits, where failures are timeouts at the episode horizon. Limits are applied independently per axis; in our insertion setup contact is dominated by the z direction, and we did not observe adverse effects from cross-axis coupling.

We further evaluate the limits in a high-force collision test by commanding 20 N in the z -direction against a rigid stop. Performance is averaged over 20 independent trials (Fig. 8). Despite measurement noise and signal smoothing in the hierarchical controller, the measured force reliably settles within a ± 1.0 N band around the 5.0 N target in an average of 0.303 s. Once settled, the controller demonstrates consistent stability, maintaining a steady-state mean error of 0.004 N with a standard deviation of 0.37 N. This confirms that the

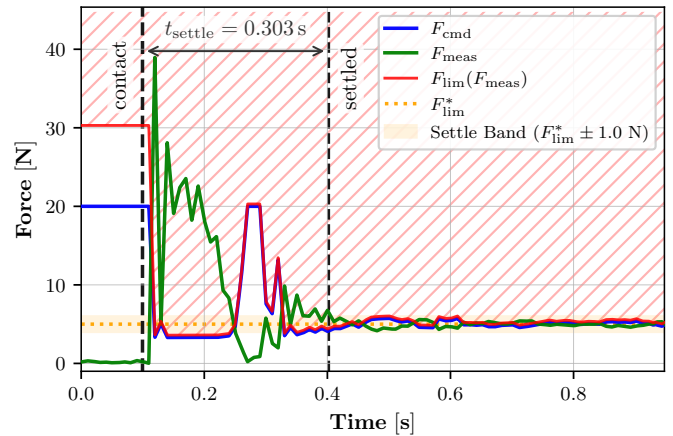


Fig. 8: **Adaptive force limits on the real system in a high-force collision test.** We command 20 N in z against a rigid stop with $F_{\max} = 30$ N and $F_{\lim}^* = 5$ N. Prior to contact, large commands are permitted; upon contact, the limit contracts and the measured force converges to F_{\lim}^* .

adaptive limits effectively mitigate impact energy and bound contact forces safely during exploration.

VI. CONCLUSION

This work introduced SHaRe-RL, a modular reinforcement learning framework showing how structured priors and interactive learning together overcome the limitations of both hand-engineered control and unstructured trial-and-error learning. Our experiments on industrial connector insertion show that this formulation improves over an unstructured human-in-the-loop RL baseline and can generalize to previously unseen connector variants. This suggests that, when guided by domain expertise already present in SMEs, RL can be both sample-efficient and safe enough for practical use, which substantially lowers the barrier to deploying learning in contact-rich industrial assembly.

A key challenge that remains is the amount of human supervision required for complex tasks. This motivates future work on more autonomous methods, such as model-based exploration [46], or self-supervised reward signals from images [47]. Our method is also orthogonal to other means of reducing real-world training time, such as offline RL pretraining [48], sim-to-real transfer [49], and the integration of large datasets and pretrained models [50].

Looking ahead, while SHaRe-RL performs well within a connector family, HMLV deployment ultimately requires broader generalization across product variants and operating conditions. A natural next step is to leverage pretrained foundation models. Recent vision-language-action (VLA) systems [51] suggest that combining large-scale pretraining, offline RL objectives, and human-in-the-loop rollouts can improve over dataset performance, yet whether such recipes transfer to specialized industrial assembly, where tasks are strongly out-of-distribution and in-domain data is limited, remains open. Identifying a robust RL fine-tuning recipe in this regime is therefore a central next step for scaling SHaRe-RL beyond narrowly scoped skills. Second, we aim to reduce the

burden of manually specifying process structure. While MP-Nets provide sample efficiency and interpretability, designing them still requires manual engineering effort. An important direction is to learn an initial, human-interpretable primitive graph from end-to-end demonstrations and iteratively refine it with human feedback, implementing a form of co-construction [52] that spans from low-level control to high-level task programming.

REFERENCES

- [1] A. Peterson, “High-Mix/Low-Volume Manufacturers Are a Sweet Spot for Collaborative Robots,” *NIST*, Jul. 2020.
- [2] Í. Elguea-Aguinaco *et al.*, “A review on reinforcement learning for contact-rich robotic manipulation tasks,” *Robotics and Computer-Integrated Manufacturing*, vol. 81, p. 102517, Jun. 2023.
- [3] M. Braun and S. Wrede, “Grey-Box Learning of Adaptive Manipulation Primitives for Robotic Assembly,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, May 2023.
- [4] A. Rajeswaran *et al.*, “Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations,” Jun. 2018.
- [5] M. Vecerik *et al.*, “Leveraging Demonstrations for Deep Reinforcement Learning on Robotics Problems with Sparse Rewards,” Oct. 2018.
- [6] J. Luo *et al.*, “Deep Reinforcement Learning for Robotic Assembly of Mixed Deformable and Rigid Objects,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2018, pp. 2062–2069.
- [7] M. and S. Wrede, “Incorporation of Expert Knowledge for Learning Robotic Assembly Tasks,” in *2020 25th IEEE International Conference on Emerging Technologies and Factory Automation*, Sep. 2020.
- [8] B. Finkemeyer, T. Kröger, and F. M. Wahl, “Executing assembly tasks specified by manipulation primitive nets,” *Advanced Robotics*, vol. 19, no. 5, pp. 591–611, Jan. 2005.
- [9] H. Beierling *et al.*, “Advancing human-robot collaboration: The impact of flexible input mechanisms,” in *Mechanisms for Mapping Human Input to Robots From Robot Learning to Shared Control/Autonomy-Workshop RSS*, vol. 2024, 2024.
- [10] J. Luo *et al.*, “Precise and Dexterous Robotic Manipulation via Human-in-the-Loop Reinforcement Learning,” Oct. 2024.
- [11] R. Cadene *et al.*, “Lerobot: State-of-the-art machine learning for real-world robotics in pytorch,” <https://github.com/huggingface/lerobot>, 2024.
- [12] M. T. Mason, “Compliance and Force Control for Computer Controlled Manipulators,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 11, no. 6, pp. 418–432, Jun. 1981.
- [13] T. Kroger, B. Finkemeyer, and F. Wahl, “A task frame formalism for practical implementations,” in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. 2004*, Apr. 2004.
- [14] D. Bargmann, W. Kraus, and M. F. Huber, “Enabling maintainability of robot programs in assembly by extracting compositions of force- and position-based robot skills from learning-from-demonstration models,” in *IROS 2024*, 2024, pp. 9227–9234.
- [15] W. Masson, P. Ranchod, and G. Konidaris, “Reinforcement Learning with Parameterized Actions,” Nov. 2015.
- [16] M. Dalal, D. Pathak, and R. Salakhutdinov, “Accelerating Robotic Reinforcement Learning via Parameterized Action Primitives.”
- [17] N. Vuong and Q.-C. Pham, “Reinforcement Learning with Parameterized Manipulation Primitives for Robotic Assembly,” Jun. 2023.
- [18] A. Padalkar *et al.*, “Learning to Close the Gap: Combining Task Frame Formalism and Reinforcement Learning for Compliant Vegetable Cutting,” in *ICINCO 2020*.
- [19] H. Hu, S. Mirchandani, and D. Sadigh, “Imitation Bootstrapped Reinforcement Learning,” May 2024.
- [20] P. J. Ball *et al.*, “Efficient Online Reinforcement Learning with Offline Data,” May 2023.
- [21] A. Nagabandi *et al.*, “Learning to Adapt in Dynamic, Real-World Environments Through Meta-Reinforcement Learning,” Feb. 2019.
- [22] P. Wu *et al.*, “DayDreamer: World Models for Physical Robot Learning,” Jun. 2022.
- [23] L. Ankile *et al.*, “From Imitation to Refinement – Residual RL for Precise Assembly,” Nov. 2024.
- [24] S. Ross, G. J. Gordon, and J. A. Bagnell, “A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning,” Mar. 2011.
- [25] M. Kelly *et al.*, “HG-Dagger: Interactive Imitation Learning with Human Experts,” Mar. 2019.
- [26] J. Luo *et al.*, “SERL: A Software Suite for Sample-Efficient Robotic Reinforcement Learning,” Feb. 2024.
- [27] P. Hartmann, J. Stranghöner, and K. Neumann, “End-to-End Low-Level Neural Control of an Industrial-Grade 6D Magnetic Levitation System,” <https://arxiv.org/abs/2509.01388v1>, Sep. 2025.
- [28] Y. Meng *et al.*, “ConBaT: Control Barrier Transformer for Safe Policy Learning,” Mar. 2023.
- [29] G. Dalal *et al.*, “Safe Exploration in Continuous Action Spaces,” Jan. 2018.
- [30] X. Wang *et al.*, “Guarding Force: Safety-Critical Compliant Control for Robot-Environment Interaction,” May 2024.
- [31] A. B. Tahmaz, R. Prakash, and J. Kober, “Impedance Primitive-augmented Hierarchical Reinforcement Learning for Sequential Tasks,” Aug. 2025.
- [32] C. C. Beltran-Hernandez, I. G. Ramirez-Alpizar, and K. Harada, “Variable Compliance Control for Robotic Peg-in-Hole Assembly: A Deep Reinforcement Learning Approach,” *Applied Sciences*, vol. 10, no. 19, p. 6923, Oct. 2020.
- [33] R. Martín-Martín *et al.*, “Variable Impedance Control in End-Effector Space: An Action Space for Reinforcement Learning in Contact-Rich Tasks,” Aug. 2019.
- [34] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, second edition ed., ser. Adaptive Computation and Machine Learning Series. Cambridge, Massachusetts: The MIT Press, 2018.
- [35] T. Haarnoja *et al.*, “Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor,” Aug. 2018.
- [36] F. Suárez-Ruiz and Q. C. Pham, “A Framework for Fine Robotic Assembly,” Sep. 2015.
- [37] A. Singh *et al.*, “End-to-End Robotic Reinforcement Learning without Reward Engineering,” May 2019.
- [38] A. Gupta *et al.*, “Reset-Free Reinforcement Learning via Multi-Task Learning: Learning Dexterous Manipulation Behaviors without Human Intervention,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, May 2021, pp. 6664–6671.
- [39] K. Neumann and J. J. Steil, “Learning robot motions with stable dynamical systems under diffeomorphic transformations,” *Robotics and Autonomous Systems*, vol. 70, pp. 1–15, Aug. 2015.
- [40] N. Hansen and X. Wang, “Generalization in Reinforcement Learning by Soft Data Augmentation,” Apr. 2021.
- [41] K. He *et al.*, “Deep Residual Learning for Image Recognition,” Dec. 2015.
- [42] S. Elfving, E. Uchibe, and K. Doya, “Sigmoid-Weighted Linear Units for Neural Network Function Approximation in Reinforcement Learning,” Nov. 2017.
- [43] G. Schoettler *et al.*, “Deep Reinforcement Learning for Industrial Insertion Tasks with Visual Inputs and Natural Rewards,” Aug. 2019.
- [44] B. Luo *et al.*, “Human-in-the-Loop Reinforcement Learning in Continuous-Action Space,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 35, no. 11, pp. 15 735–15 744, Nov. 2024.
- [45] J. Schulman *et al.*, “Proximal policy optimization algorithms,” *CoRR*, vol. abs/1707.06347, 2017.
- [46] N. Watters *et al.*, “COBRA: Data-Efficient Model-Based RL through Unsupervised Object Discovery and Curiosity-Driven Exploration,” Aug. 2019.
- [47] Y. J. Ma *et al.*, “VIP: Towards Universal Visual Reward and Representation via Value-Implicit Pre-Training,” Mar. 2023.
- [48] A. Nair *et al.*, “Learning on the Job: Self-Rewarding Offline-to-Online Finetuning for Industrial Insertion of Novel Connectors from Vision,” Feb. 2023.
- [49] L. Bergmann *et al.*, “Precision-Focused Reinforcement Learning Model for Robotic Object Pushing,” Nov. 2024.
- [50] O. X.-E. Collaboration *et al.*, “Open X-Embodiment: Robotic Learning Datasets and RT-X Models,” Jun. 2024.
- [51] P. Intelligence *et al.*, “ $\pi^*_{\{*\}}\{0.6\}$: A VLA That Learns From Experience,” Nov. 2025.
- [52] A.-L. Vollmer *et al.*, “Studying the co-construction of interaction protocols in collaborative tasks with humans,” in *4th International Conference on Development and Learning and on Epigenetic Robotics*, 2014.