

A Simulation Study of Hesitation-Aware Robotic Handoff: Safety-Efficiency Tradeoffs Across Collaborative Environment Classes

Arjun Tyagi

Data Sciences Program
Penn State Behrend
Erie, PA, USA
avt6053@psu.edu

Palash Nimesh Doshi, Uday Aditya Jain and Hriday Jagjit Matharu

Mechanical Engineering Program
Penn State Behrend
Erie, PA, USA
ppd5277@psu.edu, uaj5017@psu.edu, hjm5683@psu.edu

Abstract

Human-robot collaboration (HRC) systems commonly rely on geometric separation, nominal timing schedules, and predefined safety thresholds to manage shared-workspace interaction. While these mechanisms are essential, they do not fully account for hesitation-related human behaviors—such as pausing, delayed clearance, corrective re-entry, or rework—that can signal continued engagement and elevate overlap risk if ignored. This paper presents an environment-stratified A/B simulation study comparing a baseline robotic handoff policy (Policy A) against a hesitation-aware policy driven by trained hesitation-state inference (Policy B). Policy A follows conventional handoff timing without explicit hesitation awareness. Policy B uses a classical hesitation model, trained in Python and executed via native MATLAB inference, to trigger conservative robot-side yielding when hesitation-related states are detected. The benchmark evaluates five validated collaborative environment classes—low-conflict open workspace, narrow assembly bench, precision insertion, inspection/rework, and shared-bin access—spanning low- and high-conflict shared-zone conditions. Results show that Policy B reduces mean overlap-risk events from 0.4660 to 0.0175 per run in high-conflict environments (approximately 96.2% reduction) and from 0.1100 to 0.0080 in low-conflict environments, with task completion time increasing by approximately 1.21 and 1.00 seconds, respectively. The safety improvement is achieved primarily through increased robot hold behavior rather than through increased human wait time. These findings suggest that human hesitation can serve as a useful coordination signal in collaborative robot control, particularly in high-conflict shared workspaces where timing uncertainty and rework are common. All results are simulation-based; physical robot validation remains a direction for future work.

Keywords

Human-robot collaboration, hesitation-aware control, robotic handoff, shared workspace and A/B simulation.

1. Introduction

Collaborative robots are increasingly deployed in industrial and laboratory environments where humans and robots operate in close physical proximity (Ajoudani et al. 2018; Li et al. 2024). Applications include shared-workspace assembly, inspection, kitting, material transfer, and handoff coordination. In these settings, safe and efficient interaction depends not only on robot path planning and speed control, but also on the timing and behavioral fluency

of human–robot coordination (Lasota et al. 2017). A robot that operates on a fixed nominal schedule may perform correctly under ideal conditions but introduce risk when a human operator pauses, re-enters the shared zone, corrects a part, or otherwise deviates from the expected task sequence.

Current collaborative robot safety practice emphasizes physical separation, monitored speed reduction, protective stops, and compliance with standards such as ISO/TS 15066 (ISO 2016). These mechanisms are foundational and necessary, but they are insufficient to address subtler coordination problems in which a human is physically nearby and behaviorally uncertain (Lasota et al. 2017; Ajoudani et al. 2018; Valori et al. 2021). In shared-workspace handoff tasks, a pause is not simply a delay: it may indicate ongoing inspection, corrective action, uncertainty, or incomplete task progress. Treating hesitation as noise or delay risks premature robot motion and increased overlap-risk events—situations in which both human and robot contend for the same workspace region simultaneously.

This paper investigates hesitation-aware robotic handoff through a controlled A/B simulation study (Huck et al. 2022). The study compares a baseline handoff policy with a model-aware policy that modifies robot behavior in response to predicted hesitation-related states. Rather than asserting universal benefit from hesitation awareness, this work addresses a more specific and conditionally answerable question: under what collaborative environment conditions does accounting for human hesitation improve safety–performance tradeoffs?

The study focuses on five validated environment classes representing realistic human–robot collaboration settings with different levels of shared-zone conflict, rework likelihood, timing uncertainty, and overlap sensitivity. Results show that hesitation-aware control is most valuable in high-conflict environments, where robot and human frequently contend for shared workspace access. The paper contributes a reproducible simulation benchmark, a trained hesitation-state inference mechanism, an environment-stratified evaluation protocol, and a conditional interpretation of hesitation-aware control that honestly accounts for both safety gains and efficiency costs.

1.1 Objectives and Contributions

This paper makes the following contributions to the field of collaborative robotic control and simulation-based HRC evaluation:

A controlled A/B simulation framework comparing a baseline and a hesitation-aware robotic handoff policy across environment-stratified collaborative workspace classes.

A trained hesitation-state inference mechanism derived from a classical model, exported from Python as a structured artifact, and executed natively within MATLAB/Octave simulation without heuristic fallback.

Quantitative evidence that hesitation-aware control reduces overlap-risk events substantially in high-conflict shared workspaces, with only modest task-completion-time penalties and negligible increases in human wait time.

An honest characterization of the safety–efficiency tradeoff, including the finding that Policy B operates primarily as a conservative yielding mechanism, increasing robot hold counts rather than eliminating overlap risk through direct trajectory modification.

A conditional claim: hesitation-aware control is most beneficial in high-conflict shared-workspace environments where timing uncertainty and rework behavior are common, not universally across all collaboration settings.

1.2 Study Objectives

The specific, measurable objectives of this study are as follows:

To design and implement an A/B simulation framework in MATLAB/Octave that evaluates robotic handoff policies under consistent, reproducible conditions across multiple collaborative environment classes.

To integrate a trained classical hesitation model—derived from Python training, exported as a structured artifact, and executed via native MATLAB inference—as the decision driver for the hesitation-aware policy (Policy B).

To evaluate both policies across five validated collaborative environment classes, stratified by shared-zone conflict level (low and high), and to compare results on four quantitative metrics: overlap-risk event count, task completion time, robot hold count, and human wait time.

To quantify the safety–efficiency tradeoff between baseline and hesitation-aware control, and to assess whether the hesitation-aware policy achieves safety improvements without disproportionately increasing task time or human wait time.

To determine whether hesitation-aware control provides its greatest benefit in high-conflict shared-workspace environments, and to characterize the conditions under which hesitation awareness adds measurable value.

2. Literature Review

Human–robot collaboration research spans several interconnected domains. The present study draws on work in safety standards, human-aware robot control and intent prediction, robotic handoff and shared-zone coordination, simulation-based validation, and HRC datasets and benchmarking. This section reviews the state of each area and situates the present study within that context.

2.1 Safety Standards and Collaborative Robot Deployment

Industrial HRC safety is governed by standards that define the physical and operational limits of human–robot proximity. ISO/TS 15066:2016 provides technical guidance for collaborative robot operation, specifying human contact limits, speed and separation monitoring protocols, and risk-reduction requirements (ISO 2016). Li et al. (2024) reviewed safe HRC in industrial settings, noting that practical systems must balance safety, productivity, sensing capability, and adaptive control. While standards are essential for deployment, they are not sufficient to capture the behavioral uncertainty introduced by human hesitation, rework, or non-nominal task sequences—the focus of the present study. Broader surveys of safe HRC methods confirm that the field has consistently prioritized physical separation and speed constraints over behavioral inference (Lasota et al. 2017; Ajoudani et al. 2018). Valori et al. (2021) reinforced this point by examining the gap between validation approaches in existing standards and the broader range of safety challenges that arise in real collaborative environments.

2.2 Human-Aware Robot Control and Intent Recognition

Effective collaborative robot control requires the robot to maintain awareness of human state and intent, not just human position. Himmelsbach et al. (2021) examined human–machine differentiation in speed and separation monitoring, showing that accurately classifying human proximity and behavior affects HRC efficiency and safety outcomes. Kekana et al. (2025) reviewed intent recognition frameworks in industrial collaborative robotics, identifying trends toward model-driven inference rather than reactive position-based control. The present study extends this direction by treating hesitation specifically as a coordination signal, using trained hesitation-state inference to predict when the robot should yield rather than proceed. This is distinct from general intent recognition in that it targets a specific behavioral state—hesitation—as a safety-relevant feature. Related work on robot adaptation to human fatigue (Peternel et al. 2018) demonstrates that inferring human physiological and behavioral state can meaningfully improve coordination outcomes in co-manipulation contexts.

2.3 Robotic Handoff and Shared-Zone Coordination

Handoff tasks place particular demands on human–robot timing coordination because both parties must access shared workspace in sequence, and the sequence may be disrupted by hesitation, rework, or timing mismatches. Iori et al. (2023) studied reactive robot-to-human handover in perturbed scenarios, demonstrating that adaptive behavior during disrupted handovers improves interaction outcomes. Khanna et al. (2023) contributed a multimodal handover dataset with design implications for human–robot handovers, emphasizing that handover interaction requires both motion modeling and contextual understanding. The present study focuses on a specific failure mode in handoff: premature robot entry into a shared zone when the human has not yet completed clearance. Hesitation-aware control is proposed as a mechanism to detect this failure mode before it produces an overlap-risk event. Work on multi-level human-robot interaction strategies (Bdiwi et al. 2017) and learning collaborative behaviors from human demonstrations (Rozo et al. 2016) further supports the view that robot policies must adapt to human behavioral state rather than relying solely on fixed spatial thresholds.

2.4 Simulation-Based Validation in HRC

Simulation is an important tool for evaluating HRC policy behavior under controlled, repeatable, and varied conditions. Huck et al. (2022) used simulation to test robot system safety by generating hazardous human worker behaviors that would be impractical to reproduce reliably in hardware experiments. Valori et al. (2021) further

discussed the role of simulation-based approaches in validating safety behaviors that are difficult to evaluate through physical testing alone. The benchmark in the present study follows this precedent: MATLAB/Octave simulation is used to evaluate policy behavior across environment classes with different conflict levels, enabling reproducible comparison without the constraints of physical robot deployment. Simulation-only evaluation is a recognized limitation, and physical validation is identified as a necessary direction for future work.

2.5 HRC Datasets and Environment Benchmarking

Representative datasets and benchmark environments improve the ecological validity of HRC research. CHICO provides industrial pose forecasting data for human–robot collaboration scenarios (Sampieri et al. 2022). HA-ViD offers multi-view assembly video data annotated for actions, tools, objects, pauses, errors, and collaboration events (Zheng et al. 2023). Bello et al. (2026) introduced OpenMarcie, a multimodal dataset for action recognition in industrial environments. These datasets motivate the environment-stratified approach used in the present study, in which environment classes are designed to reflect realistic collaborative workspace conditions with varying degrees of shared-zone conflict, rework likelihood, and timing uncertainty. The benchmark environments are simulation-defined rather than derived directly from these datasets, but the environment classifications are informed by the task types and collaboration patterns they represent.

3. Methods

3.1 Research Design

This study uses a controlled A/B simulation design. Policy A is the baseline handoff policy: it follows nominal handoff behavior, does not monitor hesitation-related human states, and does not modify robot timing in response to human behavioral cues. Policy B is the hesitation-aware model policy: it uses trained hesitation-state inference to detect when the human may be hesitating, re-entering the shared zone, or completing corrective or rework actions, and in response triggers conservative robot-side interventions such as holding position or delaying shared-zone entry.

The purpose of the comparison is to determine (a) whether Policy B reduces overlap-risk events relative to Policy A, (b) whether any safety improvement is achieved with an acceptable efficiency cost, and (c) whether the benefit is environment-dependent rather than universal.

3.2 Simulation Architecture

The simulation framework was implemented in MATLAB/Octave (Huck et al. 2022; Valori et al. 2021). Each simulated environment defines a shared workspace with configurable parameters governing shared-zone geometry, human hesitation likelihood, rework probability, timing uncertainty, and overlap sensitivity. The simulation loop models human task execution, robot handoff motion, shared-zone entry decisions, and policy-specific interventions. At each decision step, the simulation records: task completion time, overlap-risk event count, robot hold count, human wait time, unnecessary slowdown count, and source-usage statistics for hesitation inference. Source-usage logging confirms whether Policy B invokes the trained model path (native MATLAB inference) or falls back to a heuristic alternative. A source-usage summary of 28,521 native-model predictions and zero fallback invocations was recorded, confirming that Policy B was exclusively model-driven.

3.3 Hesitation-State Model

The hesitation-aware policy is driven by a trained classical hesitation model. The model was trained in Python using a hesitation modeling package and exported as a structured JSON artifact. Native MATLAB inference routines were implemented to load and execute the model within the simulation loop, avoiding repeated Python calls and ensuring that the hesitation model is evaluated efficiently at each decision step.

The hesitation-state framework defines six behavioral states: normal progress, mild hesitation, strong hesitation, correction/rework, ready-for-robot action, and overlap risk. Policy B uses these predicted states to determine whether to proceed with nominal robot motion or to trigger a hold or delay. The feature inputs to the model include recent interaction timing, hesitation duration, zone occupancy signals, and correction event indicators.

3.4 Valid Benchmark Environments

The primary analysis uses five validated environment classes. These environments were designed to represent realistic HRC settings with different levels of shared-zone conflict and hesitation-relevant behavior. One invalid

stress-test condition from earlier development was excluded from the primary analysis because corrupted model behavior produced invalid comparison data; this exclusion is reported transparently and is not treated as selective result filtering.

Table 1. Valid benchmark environments and their real-world interpretations

Environment	Conflict Class	Real-World Interpretation
low_conflict_open	Low conflict	An open workspace with lower shared-zone contention and reduced probability of overlap. Hesitation-aware control provides smaller benefit in this setting.
narrow_assembly_bench	High conflict	A constrained bench where human and robot share a central workspace. Limited lateral space increases the cost of premature robot entry.
precision_insertion	High conflict	Tasks requiring alignment, micro-corrections, or insertion before handoff can proceed. Human timing is inherently variable.
inspection_rework	High conflict	A workcell where the human may inspect and return to correct a part before the robot can safely proceed. Rework loops increase hesitation frequency.
shared_bin_access	High conflict	Shared access to a bin, fixture, or transfer area used by both parties. Contention arises from concurrent access patterns.

3.5 Evaluation Metrics

Four quantitative metrics are used to evaluate and compare the policies:

Overlap-risk event count (primary safety metric): the number of simulation steps in which human and robot simultaneously occupy the shared zone, creating a potential contact or interference situation.

Task completion time in seconds (primary efficiency metric): the total time required to complete the handoff task from initiation to confirmed completion.

Robot hold count: the number of times the robot suspends or delays motion in response to a hesitation signal. This metric characterizes the mechanism and cost of the safety improvement.

Human wait time in seconds: the total time the human operator spends waiting for the robot to complete a hold or yield action. This metric assesses whether the safety improvement is achieved at the expense of the human operator. The evaluation emphasizes the tradeoff between overlap-risk reduction and efficiency penalty. A policy that reduces overlap events by increasing human wait time substantially would be a less favorable outcome than one that redistributes the delay to robot-side holds.

4. Data Collection and Benchmark Procedure

The benchmark was executed across the five validated environments under both Policy A and Policy B. For each environment, multiple simulation runs were completed and averaged to produce stable mean estimates of each metric. Each environment was evaluated across 50 independent simulation runs with randomized seeds, yielding 250 total observations per policy (50 runs $\times$ 5 environments). Seeds were drawn without replacement from a uniform random pool to ensure reproducibility while avoiding structured bias. The final benchmark outputs included environment-stratified safety, efficiency, intervention, and wait-time results, as well as CSV-format summary tables and graphical outputs.

The model-aware policy was verified through source-usage logging embedded in the simulation. The final source-usage summary recorded 28,521 native-model predictions and zero fallback invocations (Table 2), confirming that Policy B was driven exclusively by the trained hesitation model rather than by heuristic fallback behavior. This verification step is essential for interpreting Policy B results as model-driven rather than rule-driven outcomes.

The invalid stress-test environment excluded from the primary analysis generated corrupted output due to model-boundary conditions outside the training distribution. It is recorded in the benchmark log but excluded from the environment-stratified summary tables and graphical comparisons. Its exclusion is documented to support reproducibility.

5. Results and Discussion

5.1 Numerical Results

Table 3 summarizes the environment-stratified benchmark results by conflict level. In the low-conflict group (low_conflict_open), mean overlap-risk events decreased from 0.1100 under Policy A to 0.0080 under Policy B. In the high-conflict group (narrow_assembly_bench, precision_insertion, inspection_rework, shared_bin_access), mean overlap-risk events decreased from 0.4660 to 0.0175, representing an approximate 96.2% reduction. Task completion time increased by approximately 1.00 second in the low-conflict environment and 1.21 seconds in the high-conflict group. Robot hold counts increased substantially under Policy B in both conflict classes, confirming that conservative yielding is the primary mechanism of safety improvement.

Table 2. Source-usage validation for the model-aware policy (Policy B)

Native Model Predictions	Error Fallback	Stub Fallback	Script Passthrough
28,521	0	0	0

Table 3. Environment-stratified A/B benchmark summary

Conflict Level	Metric	Policy A Mean	Policy B Mean	Change	Interpretation
Low Conflict	Task completion time (s)	54.68	55.68	+1.00	Small time penalty
Low Conflict	Overlap-risk events	0.1100	0.0080	-0.1020	Safety improvement
Low Conflict	Robot hold count	10.00	521.77	+511.77	Conservative yielding
High Conflict	Task completion time (s)	73.05	74.25	+1.21	Small time penalty
High Conflict	Overlap-risk events	0.4660	0.0175	-0.4485	Substantial safety gain
High Conflict	Robot hold count	10.00	719.59	+709.59	Conservative yielding

The magnitude of the robot hold count increase—from 10 holds to more than 500 in the low-conflict case and more than 700 in the high-conflict case—indicates that Policy B is highly interventionist. This is not a failure of design; rather, it reflects the conservative nature of hesitation-aware control. The policy is calibrated to avoid overlap risk at the cost of increased holds. Whether this tradeoff is operationally acceptable depends on the specific task, throughput requirements, and tolerance for robot-side delay. These considerations are discussed further in the proposed improvements subsection.

5.2 Graphical Results (Figure 1- Figure 4)




Figure 1. Safety comparison: total overlap-risk events by environment. Policy B reduces overlap events across all five environments, with the largest absolute reductions in inspection_rework, narrow_assembly_bench, and shared_bin_access.




Figure 2. Efficiency comparison: mean task completion time by environment. Policy B is consistently slower than Policy A, but the time penalty is modest in all environments.




Figure 3. Robot hold count comparison by environment. Policy B produces substantially more robot holds, confirming that conservative yielding is the safety mechanism underlying the overlap-risk reduction.




Figure 4. Safety–efficiency tradeoff across all benchmark runs. Policy B points shift toward lower overlap-risk counts while maintaining similar task-completion-time ranges, supporting the interpretation that hesitation-aware control improves safety with modest efficiency cost.

5.3 Discussion: Why High-Conflict Environments Benefit Most

The environment-stratified results reveal a consistent pattern: hesitation-aware control provides the largest absolute safety improvements in high-conflict shared-workspace environments. This pattern can be attributed to several structural features of these environments.

In the *narrow_assembly_bench* environment, limited physical space means that both human and robot must access the same central zone, and there is little margin for timing error. When the human hesitates before clearing the shared zone, a robot operating on nominal timing has a high probability of entering while the human is still present. A policy that delays entry when hesitation is detected directly addresses this risk.

In the *inspection_rework* environment, the human is likely to return to the shared zone after an initial exit—for example, to reinspect a component or correct a fit. Baseline policies that assume one-directional human clearance are poorly suited to this pattern. Hesitation-aware control, by treating corrective re-entry as a detectable signal, reduces the probability that the robot proceeds at the moment the human returns.

In the *precision_insertion* and *shared_bin_access* environments, human timing is inherently variable because part alignment and bin access involve fine motor actions that do not complete on a fixed schedule. Nominal timing assumptions are particularly fragile in these settings, making hesitation-aware delay decisions more frequently beneficial.

In contrast, the *low_conflict_open* environment exhibits fewer overlap events under either policy because shared-zone contention is less frequent and timing uncertainty is lower. The hesitation-aware policy still reduces overlap events in this setting, but the absolute improvement is smaller, and the robot hold cost may be disproportionate to the benefit in some operational scenarios.

5.4 Proposed Improvements

The results indicate that hesitation-aware control is effective but conservative. The large increase in robot hold count suggests that Policy B may intervene more frequently than is necessary in some conditions. The following improvements are recommended for future development:

Environment-specific intervention thresholds: configure separate hesitation-detection sensitivity levels for high- and low-conflict environments so that conservative yielding is applied most aggressively where overlap risk is highest.

Confidence-aware policy scaling: reduce robot hold duration or frequency when the hesitation model reports low prediction confidence, and increase intervention intensity when confidence is high.

Cost-sensitive control: introduce a cost function that jointly optimizes overlap-risk reduction and task-time penalty, allowing the policy to select interventions proportional to the estimated risk level.

Human wait time constraint: add an explicit constraint that limits added human wait time, preventing the policy from improving robot-side safety at a disproportionate cost to the human operator.

Physical calibration: calibrate hesitation detection thresholds and model parameters using data from real human–robot handoff experiments, rather than relying solely on simulation-derived training.

5.5 Validation

The benchmark was validated in three ways. First, source-usage logging confirmed that Policy B was model-driven throughout the benchmark. The recorded summary of 28,521 native-model predictions and zero fallback invocations (Table 2) confirms that the hesitation-aware policy operated exclusively on trained hesitation-state inference rather than heuristic rules. This is essential for interpreting Policy B results as reflective of model behavior rather than rule-based approximation.

Second, the A/B comparison structure ensured a clean policy comparison. Policy A remained identical across all runs as the baseline, and Policy B differed only in the addition of hesitation-state inference and conservative yielding behavior. No other simulation parameters were changed between policies.

Third, the benchmark was restricted to five validated environments. The one invalid stress-test condition excluded from the primary analysis is documented transparently. Its exclusion improves the interpretability of the final results and is not treated as selective removal of unfavorable data.

These validation steps support a conditional interpretation of the results: hesitation-aware control reduces overlap-risk events in the evaluated environments, and the mechanism is model-driven conservative yielding. The generalizability of the results beyond the simulated environments and the trained model distribution requires physical validation.

Statistical hypothesis testing using the Wilcoxon signed-rank test confirms that all observed differences between Policy A and Policy B are statistically significant. For overlap-risk event count, the test yields $p = 1.90 \times 10^{-22}$, confirming that the safety improvement is not attributable to chance (effect size $d = 1.06$, large). For task completion time, $p = 2.99 \times 10^{-41}$ with a small effect size ($d = -0.10$), consistent with the modest but real efficiency cost. Robot hold count differences are highly significant ($p = 9.22 \times 10^{-43}$, $d = -5.75$), confirming that conservative yielding is the dominant mechanism. Human wait time differences, while statistically significant ($p = 3.61 \times 10^{-29}$), carry a negligible effect size ($d = -0.02$), confirming that the policy does not materially burden the human operator. All tests were conducted at $\alpha = 0.05$ across $n = 250$ observations per policy.

5.6 Limitations

This study is subject to several limitations that should be considered when interpreting the results and assessing the generalizability of the findings.

Simulation-only validation: all results are derived from a MATLAB/Octave simulation framework. No physical robot experiments were conducted. Simulation enables repeatable, controlled policy comparison but may not capture all dynamics of real human–robot interaction, including physical contact forces, sensing noise, actuation latency, and unpredictable human behavior.

Synthetic hesitation training: the hesitation model was trained using a Python hesitation modeling package rather than labeled data from real human–robot handoff experiments. The ecological validity of the model's hesitation-state predictions in physical deployments has not been established.

No physical robot deployment: the hesitation-aware policy has not been implemented or tested on any physical collaborative robot platform. Claims about operational benefit are conditional on simulation results only.

Sensitivity to model calibration: the policy behavior of Policy B is sensitive to the threshold and weighting parameters of the hesitation model. The results reported here reflect a single model configuration; sensitivity analysis across parameter ranges was not conducted in this study.

Excluded invalid environment: one stress-test benchmark condition was excluded from the primary analysis due to invalid model behavior at distribution boundaries. Although this exclusion is documented, its presence in the original benchmark design represents a limitation in the robustness of the testing protocol.

Conservative intervention behavior: the large robot hold counts produced by Policy B suggest that the current policy is more conservative than may be operationally appropriate in low-risk conditions. A more adaptive policy that scales intervention to estimated conflict level would be preferable in practice.

Limited environment diversity: the benchmark covers five environment classes. Additional environment types—including dynamic human trajectories, multi-person workspaces, or sensor-occlusion conditions—were not included in the primary validated analysis.

6. Conclusion

This paper presented an environment-stratified A/B simulation study comparing a baseline robotic handoff policy (Policy A) with a hesitation-aware model-driven policy (Policy B) across five validated collaborative workspace environments. The hesitation-aware policy used trained hesitation-state inference, implemented via native MATLAB model execution, to trigger conservative robot-side yielding when hesitation-related behavioral states were predicted.

Across all five environments, Policy B substantially reduced overlap-risk events relative to Policy A. The largest absolute safety improvements occurred in high-conflict shared-workspace environments—narrow assembly bench, precision insertion, inspection/rework, and shared-bin access—where shared-zone contention, rework probability, and timing uncertainty are highest. In these environments, mean overlap-risk events fell from 0.4660 to 0.0175 per run, a reduction of approximately 96.2%, with a task completion time penalty of approximately 1.21 seconds. Human wait time remained nearly unchanged, indicating that the safety improvement was achieved primarily through increased robot hold behavior rather than through added human delay.

The main conclusion is that human hesitation should not be treated merely as delay or noise in shared-workspace collaboration (Lasota et al. 2017; Ajoudani et al. 2018). In environments where timing uncertainty and rework are common, hesitation can serve as a meaningful coordination signal (Kekana et al. 2025). Incorporating that signal into robot control—through trained inference rather than heuristic rules—can reduce overlap-risk events while preserving task efficiency within acceptable bounds.

Future work should address the limitations of this study through several directions: physical robot deployment and validation of the hesitation-aware policy on a collaborative robot platform; calibration of the hesitation model using real handoff data; environment-specific sensitivity tuning to reduce unnecessary robot holds in low-conflict settings; confidence-aware policy scaling; and extension of the benchmark to cover a broader range of collaborative workspace conditions including multi-person environments, dynamic human trajectories, and sensor-occlusion scenarios.

References

Bello, H., Ray, L. S. S., Sorysz, J., Suh, S. and Lukowicz, P., OpenMarcie: Dataset for multimodal action recognition in industrial environments, arXiv preprint arXiv:2603.02390, 2026.

- Himmelsbach, U. B., Wendt, T. M., Hangst, N., Gawron, P. and Stiglmeier, L., Human-machine differentiation in speed and separation monitoring for improved efficiency in human-robot collaboration, *Sensors*, vol. 21, no. 21, article 7144, 2021.
- Huck, T. P., Ledermann, C. and Kroger, T., Testing robot system safety by creating hazardous human worker behavior in simulation, *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 770–777, 2022.
- Iori, F., Perovic, G., Cini, F., Mazzeo, A., Falotico, E. and Controzzi, M., DMP-based reactive robot-to-human handover in perturbed scenarios, *International Journal of Social Robotics*, vol. 15, pp. 233–248, 2023.
- ISO, ISO/TS 15066:2016 Robots and robotic devices — Collaborative robots, International Organization for Standardization, Geneva, Switzerland, 2016.
- Kekana, M., Du, S., Steyn, N., Benali, A. and Djerroud, H., A review of human intention recognition frameworks in industrial collaborative robotics, *Robotics*, vol. 14, no. 12, article 174, 2025.
- Khanna, P., Bjorkman, M. and Smith, C., A multimodal data set of human handovers with design implications for human-robot handovers, *Proceedings of the 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pp. 1843–1850, Busan, Korea, 2023.
- Li, W., Hu, Y., Zhou, Y. and Pham, D. T., Safe human-robot collaboration for industrial settings: a survey, *Journal of Intelligent Manufacturing*, vol. 35, pp. 2235–2261, 2024.
- Sampieri, A., D'Amely di Melendugno, G. M., Avogaro, A., Cunico, F., Setti, F., Skenderi, G., Cristani, M. and Galasso, F., Pose forecasting in industrial human-robot collaboration, *Computer Vision — ECCV 2022, Lecture Notes in Computer Science*, vol. 13698, pp. 51–69, 2022.
- Valori, M., Scibilia, A., Fassi, I., Saenz, J., Behrens, R., Herbster, S., Bidard, C., Lucet, E., Magisson, A. and Schaake, L. et al., Validating safety in human-robot collaboration: standards and new perspectives, *Robotics*, vol. 10, no. 2, article 65, 2021.
- Zheng, H., Lee, R. and Lu, Y., HA-ViD: A human assembly video dataset for comprehensive assembly knowledge understanding, *Advances in Neural Information Processing Systems 36, Datasets and Benchmarks Track*, 2023.
- Ajoudani, A., Zanchettin, A. M., Ivaldi, S., Albu-Schäffer, A., Kosuge, K. and Khatib, O., Progress and prospects of the human-robot collaboration, *Autonomous Robots*, vol. 42, no. 5, pp. 957–975, 2018.
- Bdiwi, M., Pfeifer, M. and Sterzing, A., A new strategy for ensuring human safety during various levels of interaction with industrial robots, *CIRP Annals*, vol. 66, no. 1, pp. 453–456, 2017.
- Lasota, P. A., Fong, T. and Shah, J. A., A survey of methods for safe human-robot interaction, *Foundations and Trends in Robotics*, vol. 5, no. 4, pp. 261–354, 2017.
- Peternel, L., Tsagarakis, N., Caldwell, D. and Ajoudani, A., Robot adaptation to human physical fatigue in human-robot co-manipulation, *Autonomous Robots*, vol. 42, no. 5, pp. 1011–1021, 2018.
- Rozo, L., Calinon, S., Caldwell, D. G., Jimenez, P. and Torras, C., Learning physical collaborative robot behaviors from human demonstrations, *IEEE Transactions on Robotics*, vol. 32, no. 3, pp. 513–527, 2016.

Biographies

Arjun Tyagi is an undergraduate student in the Data Sciences Program at Penn State Behrend, expected to graduate in 2028. His academic interests include artificial intelligence, machine learning, autonomous systems, simulation-based benchmarking, and applied human-robot collaboration. In this project, he served as the student PI/coordinator while contributing to the data-driven hesitation modeling, simulation framing, and manuscript development.

Palash Nimesh Doshi is an undergraduate Mechanical Engineering student at Penn State Behrend, expected to graduate in 2029. His research interests include human-robot collaboration, mechanical systems, modeling human hesitation, and designing shared industrial systems that are safer and more intuitive. He is driven by curiosity, machine design, innovation, and practical engineering applications.

Uday Aditya Jain is a first-year Mechanical Engineering student at Penn State Behrend, expected to graduate in 2029. He has hands-on experience in MATLAB/Simulink vehicle dynamics simulation and is proficient in SolidWorks and Python. His interests include mechanical simulation, design, robotics, and the integration of computational tools into engineering workflows.

Hriday Jagjit Matharu is an undergraduate Mechanical Engineering student at Penn State Behrend, expected to graduate in 2028. His interests include mechanical design, simulation-based engineering, collaborative robotics, and safety-focused analysis of human-machine interaction. He contributed to the mechanical engineering interpretation and collaborative workspace framing of the study.