

Towards Fast Reprogramming and Adaptive Reproduction of Contact-rich Assembly

Dimitrios Rakovitis¹, Vamsi Krishna Origanti¹, Vinzenz Bargsten¹, Adrian Danzglock¹

¹All authors are with Robotics Innovation Center of German Research Center for Artificial Intelligence GmbH (DFKI), Bremen, Germany.

Correspondence*:
Dimitrios Rakovitis
dimitrios.rakovitis@dfki.de

2 ABSTRACT

3 Modern manufacturing demands flexible, robust robotic assembly systems capable of handling
4 variable part geometries and dynamic task configurations. Current approaches often suffer from
5 limited generalization, high sample complexity, and the need for extensive reconfiguration or
6 retraining when task parameters change. This paper addresses these limitations by introducing
7 a unified framework for rapid, intuitive programming and adaptive reproduction of contact-
8 rich assembly using only force/torque and proprioceptive sensing. The approach combines
9 synchronized wrench–motion Dynamic Movement Primitives (wDMPs), an uncertainty-aware
10 Adaptive Model Predictive Controller (AMPC), and a neural contact classifier based on Adaptive
11 Resonance Theory (ART). Training is performed in two steps: (i) a single kinesthetic demonstration
12 for teaching coupled motion–wrench profiles; and (ii) an assistive reproduction in which the
13 robot collects contact data from a nominal (successful) operation to train ART for alignment
14 detection and a Gaussian Mixture Model (GMM) for uncertainty estimation. During autonomous
15 reproduction, the uncertainty continuously adapts the MPC’s model to enable compliant contact
16 handling, while the ART classifier detects misalignments and coordinates transitions between
17 assembly stages. Evaluated on standard benchmarks and real-world industrial scenarios,
18 including peg-in-hole, plug-insertion, and disc brake assemblies, the framework demonstrates
19 improved robustness, and transferability over baseline approaches. This contributes toward
20 enabling agile, customizable production with minimal reprogramming effort.

1 INTRODUCTION

21 Modern manufacturing increasingly demands the ability to produce a diverse range of products, requiring
22 assembly systems to handle frequent changes in part geometry, alignment, and sequencing. This variability
23 challenges traditional robotic assembly solutions, which are typically finely engineered or trained for
24 well-defined tasks and variations (Luo et al., 2019; Tang et al., 2023; Noseworthy et al., 2025; Guo et al.,
25 2025; Wu et al., 2025; Yan et al., 2021; Jha et al., 2022; Schoettler et al., 2020; Goyal et al., 2024; Kim
26 et al., 2020; Chang et al., 2022; Morgan et al., 2023). Even then, such solutions are often labor-intensive

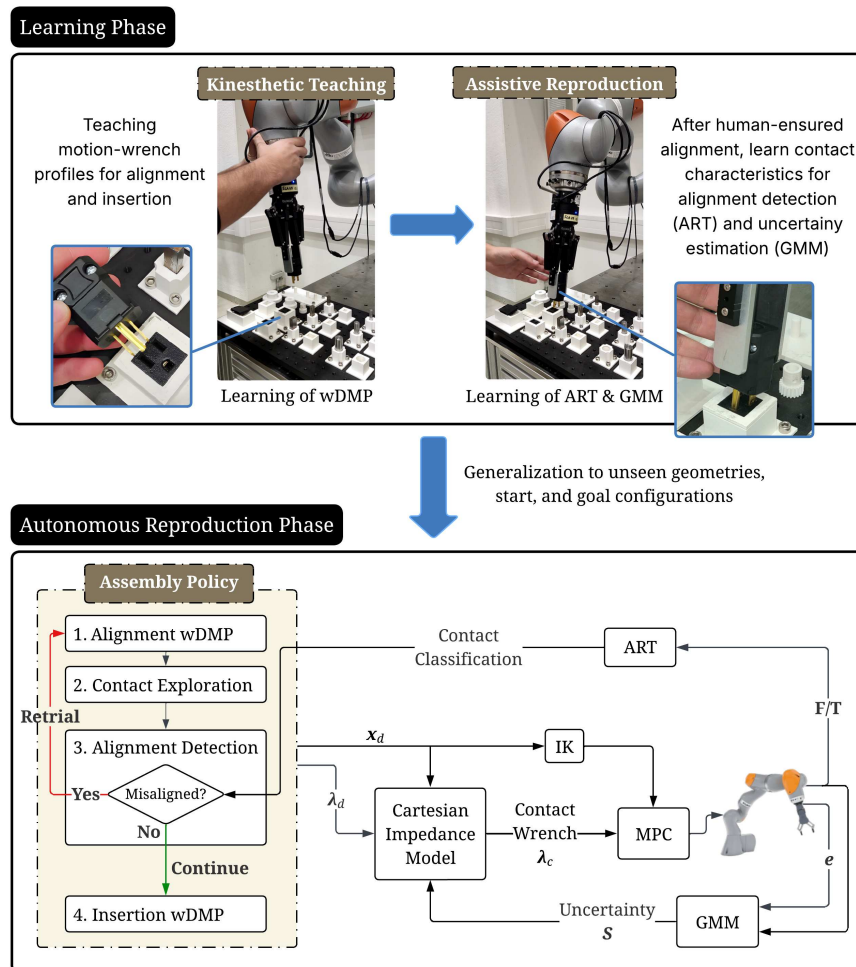


Figure 1. Graphical abstract of the proposed framework.

27 and time-consuming, which underscores the need for systems capable of robust, adaptive behavior with
 28 minimal manual reprogramming.

29 To address these challenges, we propose a novel framework that enables fast, intuitive robot programming
 30 and adaptive reproduction of previously unseen, contact-rich assembly tasks using only a force/torque
 31 (F/T) sensor and error measurements. The method integrates Dynamic Movement Primitives (DMPs),
 32 Adaptive Model predictive Control (AMPC), and an Adaptive Resonance Theory (ART)-based contact
 33 classifier. DMPs encode human-taught assembly demonstrations as stable nonlinear dynamical systems,
 34 enabling the recorded motions and wrenches to generalize smoothly to new start/goal states and fixtures
 35 (Ijspeert et al., 2013). AMPC is a model-based predictive controller in which the robot's prediction model
 36 is updated online from data (Rakovitis and Mronga, 2024), allowing adaptive and compliant resolution
 37 of errors during reproduction of the learned assembly, e.g., invalid contacts due to misalignments. An
 38 ART-based contact classifier is a neural network (NN) that incrementally classifies F/T patterns under a
 39 vigilance criterion to recognize known contacts and flag novel ones on the fly (Bargsten et al., 2025).

40 Specifically, for each new assembly task family (e.g. peg-in-hole or plug-insertion), the approach lets a
 41 user train the system in two demonstrations: (i) a kinesthetic teaching, and (ii) an assistive reproduction.
 42 At first, motion and wrench profiles are taught from a single, nominal demonstration using synchronized
 43 wrench-motion DMPs (wDMPs). The user segments the demo into sub-wDMPs that represent sequential

44 stages, e.g. alignment and insertion. These wDMPs generalize across varying start and goal configurations
45 and supply AMPC with the motion/wrench references needed to reproduce the task, while adapting for
46 different part geometries. Secondly, the assistive reproduction takes place for the same goal position and
47 parts. Throughout, AMPC uses a Cartesian Impedance Model (CIM) to represent interaction dynamics
48 inside the MPC prediction. The demonstrator provides small corrections to align the end-effector (EE)
49 properly with the assembly structure. Once aligned, the robot performs a gentle "jiggling" motion along
50 the assembly axis to gather contact data. The resulting F/T signals train the ART-based classifier to detect
51 correct alignment and generalize this knowledge for similar parts. Then, the insertion phase takes place,
52 collecting data to fit a Gaussian Mixture Model (GMM) that captures the joint distribution of EE pose
53 errors and contact wrenches under nominal operation. This model is used to estimate uncertainty based on
54 the likelihood of new observations.

55 During autonomous reproduction, execution is governed by an assembly policy that switches between
56 four stages, while AMPC tracks references from the active stage: (i) Alignment wDMP, (ii) Contact
57 exploration, (iii) Alignment detection, (iv) Insertion wDMP. After an alignment attempt, force-driven
58 jiggings are applied at the EE about the estimated goal direction, guided by tracking errors to resolve
59 possible misalignment errors (e.g., due to modelling or goal-estimation errors). Then, during alignment
60 detection, jiggings are applied only in the assembly axis; the ART classifier monitors F/T data to flag
61 undesired contacts (e.g. misalignments or unintended collisions) and acts as a scheduler, triggering the
62 insertion wDMP upon confident alignment or initiating a retrieval when the context deviates from known
63 patterns. Throughout, the estimated GMM uncertainty modulates the AMPC compliance in real-time by
64 adjusting the parameters of the CIM (stiffness and desired wrench). This formulation is selected as it
65 enables adaptive and compliant contact exploration in case of misalignments.

66 The approach is evaluated experimentally using a 7-DOF KUKA LBR iiwa14 R820 manipulator, equipped
67 with either the Robotiq 2F-85 adaptive gripper or the OnRobot 3FG25 gripper. The OnRobot 3FG25
68 includes an integrated F/T sensor that gives measurements at 140Hz, while the Robotiq 2F-85 setup
69 incorporates an external Robotiq FT300 sensor mounted before the gripper and working at 100Hz. The
70 robot performs a wide range of contact-rich assembly tasks, starting with the standard IndustRealKit
71 benchmark (Tang et al., 2023). This includes classic peg-in-hole tasks with cylindrical and rectangular
72 pegs of varying sizes, two- and three-prong plug insertions (which require significant insertion force),
73 and multi-stage gear assemblies involving both peg insertion and gear teeth alignment. Beyond these
74 benchmarks, we evaluate the system on a real-world industrial scenario: a multi-stage disc brake assembly,
75 requiring the handling of heavy components and the application of substantial forces. Each task is repeated
76 across multiple trials to obtain statistically meaningful results. As our novelty lies primarily in the adaptive
77 reproduction of kinesthetically taught policies, the approach is compared against two state-of-the-art
78 baseline controllers (classic Cartesian impedance control (Origanti et al., 2025), and error-based predictive
79 variable impedance control (Anand et al., 2023)), and ablation baselines to highlight the importance of key
80 components during reproduction.

81 Overall, the contributions of this work are listed below:

- 82 • We propose a novel, user-friendly framework for fast programming and adaptive reproduction of
83 contact-rich, robotic assembly tasks, that integrates DMPs with uncertainty-aware AMPC, which
84 enables compliant and adaptive contact handling,
- 85 • We combine the above with a NN contact classifier based on ART, for real-time detection of undesired
86 contacts and misalignments, which also serves as a scheduler for assembly stage transitions and retrievals.

87 • We provide comprehensive experimental validation on diverse industrial and benchmark scenarios,
88 demonstrating improved handling of misalignments, and adaptability over baseline approaches.

89 The remainder of the paper is structured as follows: Section 2 details related works, Section 3 outlines
90 the proposed methodology, Section 4 discusses the experimental results, and Section 5 concludes with a
91 summary of findings and potential directions for future work.

2 RELATED WORK

92 Prior research on learning, controls, and contact classification for contact-rich tasks is detailed in the
93 following subsections.

94 2.1 Related works on learning for assembly

95 Recent research in robotic assembly has explored a wide range of learning-based approaches to address
96 the challenges of contact-rich insertion tasks. Reinforcement learning (RL) methods have shown promise
97 for learning robust, low-tolerance behaviors, e.g. with Guided policy search (Levine et al., 2015), and
98 RL-based variable impedance control (Luo et al., 2019). Simulation-based RL pipelines such as IndustReal
99 (Tang et al., 2023) and FORGE (Noseworthy et al., 2025) have demonstrated successful sim-to-real
100 transfer with tight-tolerance insertions, through domain randomization, signed-distance rewards, or force
101 conditioning. Meta-RL approaches (Schoettler et al., 2020) leverage shared structure across insertion tasks
102 to enable rapid adaptation from simulation to real-world scenarios. Similarly, SRSA (Guo et al., 2025)
103 retrieves relevant skills from a pre-existing policy library using predicted transfer success, and fine-tunes
104 them on new assembly tasks using Proximal Policy Optimization (PPO) combined with self-imitation
105 learning.

106 Supervised and imitation learning approaches have also proven effective. In (Yan et al., 2021), a
107 contact-state recognition model is trained on F/T data, collected from various inclined peg-in-hole
108 configurations. A support vector machine (SVM) with a Gaussian kernel enables accurate classification of
109 the contact state, which then informs the parameters of an adaptive impedance controller to compliantly
110 correct misalignments. In (Jha et al., 2022), DMPs are used to learn a nominal insertion trajectory
111 from human demonstrations, with a corrective compliance policy to handle vision-based goal estimation
112 errors. A generalized accommodation controller bounds contact forces for safe exploration and data
113 collection, while a Gaussian Process (GP) trained on F/T data predicts misalignments to guide successful
114 insertions despite pose errors. DMPs (Ijspeert et al., 2013) have been extensively adapted to represent
115 forces and complex, multi-modal behaviors. Several studies have incorporated force feedback into DMP
116 frameworks (Pastor et al., 2009; Hoffmann et al., 2009), while others have employed basis function
117 regression to model and reproduce force profiles (Ude et al., 2010; Kramberger et al., 2017). Extensions
118 such as Task-Parameterized DMPs (Calinon, 2016) and Probabilistic Movement Primitives (ProMPs)
119 (Paraschos et al., 2013) enhance spatial generalization and capture variability across several demonstrations.
120 (Origanti et al., 2022) showcases automatic extraction of skill parameters from human demonstrations,
121 to replicate both motion and force trajectories in simulation, particularly for contact-rich manipulation.
122 Complementarily, RVT-2 (Goyal et al., 2024) introduces a multi-view vision-based system that learns
123 high-precision manipulation from just a few demonstrations. It uses supervised behavioral cloning to map a
124 third-person RGB-D input and language instructions to key-frame poses, with a multi-view transformer and
125 a coarse-to-fine inference strategy enabling millimeter-level accuracy using only visual data. TacDiffusion
126 (Wu et al., 2025) leverages diffusion models trained on expert demonstrations using cuboid pegs to map

127 tactile observations to force-domain actions, achieving a 95.7% zero-shot transfer success rate on novel
128 peg-in-hole tasks including cylinder, prism, and key-shaped pegs.

129 Imitation learning has also been combined with RL in (Kim et al., 2020) by using a human demonstration
130 to train a NN-based movement primitive (NNMP), which constrains RL to a known trajectory manifold.
131 With a properly designed reward function, this enabled the agent to learn high-precision contact tasks
132 while minimizing applied forces. In (Chang et al., 2022), an assembly task learned via motion-force
133 DMPs is reproduced with a low-level admittance controller, whose stiffness is tuned by RL, enabling
134 real-time impedance adaptation. Other works, using compliance-enabled strategies (Morgan et al., 2023)
135 or neglecting force feedback (Park et al., 2017) offer hardware-efficient alternatives that exploit passive
136 compliance or implicit search behaviors.

137 Despite these advancements, several limitations are shared across the above works. Many approaches
138 are sample-inefficient, requiring extensive manual training or perfect simulators that can take several
139 hours to days (Luo et al., 2019; Tang et al., 2023; Noseworthy et al., 2025; Guo et al., 2025; Wu et al.,
140 2025; Yan et al., 2021; Jha et al., 2022; Calinon, 2016; Paraschos et al., 2013) or additional multiple
141 real-world trials ($\approx 10 - 30$) for transfer learning (Schoettler et al., 2020; Goyal et al., 2024; Kim et al.,
142 2020; Chang et al., 2022). Generalization to arbitrary geometries or unseen part configurations is limited,
143 often requiring timely retraining or fine re-engineering (Morgan et al., 2023). Additionally, many of the
144 approaches, consider cut-off forces or fixed compliance (Schoettler et al., 2020; Park et al., 2017; Morgan
145 et al., 2023; Noseworthy et al., 2025) for safe exploration, which may result in task failure in tasks involving
146 significantly larger forces, e.g. hard plug insertion or sliding on high friction surfaces. Adapting these
147 systems for rapidly changing contact-rich assembly tasks, typically entails significant reconfiguration. For
148 example, in a manufacturing environment this would result in costly extended production downtime (Liu
149 et al., 2012), whenever small changes in production are introduced. This highlights the need for flexible,
150 intuitive, and data-efficient methods that can rapidly adapt to diverse changes in assembly settings. For
151 these reasons, in this work we rely on DMPs to teach motion-wrench profiles from a single demo in
152 Cartesian space (Fabisch, 2024), and reproduce them via an adaptive controller to improve generalization
153 to different part geometries, start and goal configurations.

154 **2.2 Related works on Controls**

155 Typically the above works rely on compliant systems (passive or active), with many of them using
156 Cartesian Impedance Control (CIC) to introduce the required task space compliance for safe contact
157 interaction. However, CIC is inherently reactive; it lacks prediction and cannot natively enforce constraints
158 (e.g., joint/torque limits), so it often relies on heuristic task-specific safety margins (Origanti et al., 2025).

159 Recent (non-assembly) studies combine CIC with MPC into model-based impedance- (MPIC)
160 (Bednarczyk et al., 2020) or variable-impedance control (MPVIC) (Thelenberg and Ott, 2024; Anand et al.,
161 2023) for tasks involving contact uncertainties or variable stiffness objectives. By embedding impedance
162 dynamics in MPC, these approaches retain compliance while enabling look-ahead planning, explicit
163 constraint handling, and proactive adaptation across contact transitions. In particular, (Thelenberg and
164 Ott, 2024; Anand et al., 2023) leverage the impedance model inside the MPC to forecast the stiffness and
165 damping commanded by a low-level variable-impedance controller, improving task adaptability.

166 However, (Bednarczyk et al., 2020) enforces a fixed stiffness profile over the entire task, limiting the
167 controller's ability to adapt compliance dynamically based on environmental conditions or task phases.
168 Conversely, (Thelenberg and Ott, 2024; Anand et al., 2023) drive stiffness adaptation by the magnitude of
169 tracking errors, increasing stiffness when the robot is far from its target. While intuitive, this heuristic can

170 lead to unphysical or unsafe stiffness changes in contact-rich scenarios, where increased stiffness at the
171 wrong moment may cause damage, or task failure during contact transitions.

172 Other works, handle contact uncertainty by using Adaptive MPC. This method adjusts the MPC model or
173 cost/constraints in real-time based on a contact model that is either learned, e.g. via RL (Xu et al., 2022b),
174 supervised learning (Rakovitis and Mronga, 2024), or meta-learning (Saviolo et al., 2024; Anne et al.,
175 2021; Arcari et al., 2023), or predicted via adaptive control or system identification techniques (Minniti
176 et al., 2021; Xu et al., 2022a). These works have mostly been used in mobile manipulators, quadrupeds,
177 and quadrotors for compensating unknown dynamics, such as payloads and disturbances.

178 Although, all the above control approaches are achieving great results on dealing with contact
179 uncertainties, their effectiveness on fine assembly tasks have yet to be validated, where high precision,
180 delicate contact handling, and sub-millimeter accuracy are required. This highlights the need for more
181 robust studies on context-aware controllers that account for both contact dynamics and safety-critical
182 constraints in fine and diverse assembly contexts. Hence, in this work we extend MPC with a Cartesian
183 impedance contact model, whose stiffness and desired wrench are continuously adapted via uncertainty,
184 estimated from the likelihood of error/wrench measurements relative to a GMM. This enables adaptive and
185 compliant reproduction of the learned assembly tasks. We select GMMs, as they have been shown to be
186 highly sample-efficient, especially in high-dimensional spaces, and very fast to train (Calinon et al., 2007;
187 Rakovitis and Mronga, 2024). GMMs have also been used before for anomaly and out-of-distribution
188 detection (Zong et al., 2018; Iwata and Kumagai, 2022).

189 **2.3 Related works on Contact Classification**

190 The detection and classification of contacts between a robot and its environment is essential for inferring
191 the current state of the system. Prior research broadly falls into two focal points. The first addresses
192 the detection of unintended collisions, aiming to mitigate impact forces via reflexive counter-actions.
193 Classically, this type of approaches compare an estimated impact metric against a pre-defined threshold.
194 Because robotic assembly inherently involves purposeful contact, our work aligns with the second focal
195 point: the classification of intended contacts to obtain a more fine-grained assessment of the contact forces
196 when performing an assembly task.

197 Several works have tried to bridge the gap between the two focal points by jointly distinguishing
198 unintended collisions from intended interactions. For example, (Cho et al., 2012) differentiates collisions
199 from intended contact by monitoring the rate of change in joint torque measurements, while frequency-
200 domain analyses have also shown promise (Kouris et al., 2016, 2018). In robotic assembly, however, the
201 classification into only two classes (three with no-contact class) based on thresholds is insufficient to assess
202 the successful insertion and joining of rigid and elastic parts with tight tolerances. To address this, some
203 works such as (Pankert and Hutter, 2023) incorporate visual monitoring with prior knowledge to improve
204 kinematic-level performance by fusing CAD models, tactile cues, and particle simulation to refine object
205 localization.

206 In contrast, our objective is an assessment driven by contact wrenches, that does not depend on precise
207 prior models or expert supervision. Other domains underscore the potential of such signals, including
208 classification from joint torque measurements (Iskandar et al., 2024) and from acoustic vibration sensing
209 (Liu and Chen, 2024). Within robotic assembly specifically, two preliminary works have shown the
210 feasibility of employing an incremental machine learning approach for continuous classification of episodes
211 encoded by the frequency magnitudes of joint torque measurements (Bargsten and Kirchner, 2023) or EE
212 wrench measurements (Bargsten et al., 2025). In these approaches, time-series measurements are encoded

Algorithm 1: Learning of contact-rich assembly

Input : Single kinesthetic demonstration $D = \{(\mathbf{x}^{(i)}, \mathbf{q}^{(i)}, \boldsymbol{\lambda}^{(i)})\}_{i=1}^N$ in gravity mode

Output : Stage wDMPs $\{\text{wDMP}_{\text{align}}, \text{wDMP}_{\text{insert}}\}$, ART classifier, GMM $p(\mathbf{z})$

1) Kinesthetic teaching;

Learn forcing terms $\{\mathbf{f}_x, \mathbf{f}_q, \mathbf{f}_\lambda\}$ (3) of DMPs (2), (6), (9) from D ; Segment task into *alignment* and *insertion*; fit one synchronized wDMP per segment;

2) Assistive reproduction;

Reproduce the task once (both segments sequentially) with non-uncertain AMPC (3.4); wDMP provides the reference pose and wrench; human provides small corrections for EE alignment;

Once aligned: execute axial jiggling along assembly Z-axis (12); collect contact F/T data; train ART classifier;

Continue with insertion: collect error/wrench measurement pairs $\mathcal{D}^{(i)} = \{\mathbf{z}^{(i)} = [\mathbf{e}^{(i)}, \boldsymbol{\lambda}_m^{(i)}]\}_{i=1}^N$ under nominal insertion; fit GMM $p(\mathbf{z})$ (19) on \mathcal{D} ;

213 using a short-time Fourier transform (STFT), producing compact, episodic signatures that lend themselves
214 to real-time classification via Adaptive Resonance Theory (ART).

215 ART originates in cognitive science and models dynamic processes for learning and adapting short- and
216 long-term memory in the human brain (Grossberg, 1976; Carpenter and Grossberg, 1987). Building on this
217 foundation, numerous algorithmic variants (Leonardo Enzo Brito da Silva et al., 2019) have been developed
218 that simplify and operationalize these principles. ART yields match-based learning methods that rely on
219 input similarity, in contrast to error-based batch learning such as backpropagation. To this end, we adopt
220 Distributed Dual Vigilance Fuzzy ART (DDVFA) (Brito da Silva et al., 2020), a variant of ART in which
221 each learned contact-state (class) is a small Fuzzy ART (Carpenter et al., 1991) network. This nested design
222 represents classes as groups of sub-classes and naturally captures arbitrarily shaped, heterogeneous clusters.
223 Because of its dual vigilance mechanism, DDVFA naturally supports fine-grained, online adaptation
224 without catastrophic forgetting and offers sample efficiency for rare contact events. However, in our
225 framework the training of DDVFA is only enabled by the operator during the assistive reproduction demo.
226 The time-series of the EE wrench, gathered at this phase are transformed to the frequency domain (STFT)
227 and used to fit DDVFA to distinguish nominal alignment from undesired contacts. The trained classifier is
228 then frozen and deployed as a real-time context monitor and scheduler triggering insertion or a retrieval, to
229 enable adaptive stage transitions during contact-rich assembly.

3 METHODOLOGY

230 This section describes the proposed methodology in five subsections covering: (A) the DMPs formulation,
231 (B) a force-based contact exploration strategy, (C) the contact classification module via ART, (D) the MPC
232 problem, and (E) the estimation of uncertainty via a GMM and its use to adapt MPC in real-time. Our
233 framework consists of two parts: (i) the learning, and (ii) the reproduction phase. The complete pipelines
234 for each phase, are detailed in algorithms 1 and 2, respectively.

235 3.1 Wrench-Motion Dynamic Movement Primitives (wDMPs)

236 Each assembly task is specified by the parts to be assembled (whose geometry and required interaction
237 forces may be unknown), together with an initial pose and a fixtured goal pose in Cartesian space (which are
238 known). Hence, each task is encoded by two DMPs that share a single phase variable, thereby synchronizing
239 pose and wrench generation over time:

Algorithm 2: Adaptive autonomous reproduction of contact-rich assembly**Input** : Stage wDMPs $\{\text{wDMP}_{\text{align}}, \text{wDMP}_{\text{insert}}\}$, ART classifier, GMM $p(\mathbf{z})$ **Output** : Assembly execution**Global:** Task reproduced with uncertainty-aware AMPC (3.4 + 3.5); The stiffness and desired wrench (given by wDMP) of a Cartesian impedance contact model (18)) is adapted online from GMM uncertainty (21), (23), (24).**1) Execute alignment wDMP** $\text{wDMP}_{\text{align}}$.**2) Execute contact-force exploration:** force-based jiggling in all directions (10) to resolve misalignments.**3) Alignment detection:** force-based jiggling in assembly Z-axis (12); ART classifies contacts:**if** ALIGNED **then**

└ Proceed to Step 4

else└ **Retry once:** redo Steps 1-3,
└ If still NOTALIGNED: **abort****4) Execute insertion wDMP** $\text{wDMP}_{\text{insert}}$.

- 240 1. a **Cartesian DMP** that generates EE pose trajectories in $SE(3)$ (3D position and 4D quaternion),
241 2. a **Wrench DMP** that reproduces the EE wrench (3D force and 3D torque).

242 We refer to this pair as a *wrench–motion DMP* (*wDMP*). By jointly learning motion and the associated
243 interaction wrenches, wDMPs extend standard DMPs from pure kinematics to coupled motion-wrench
244 behavior, enabling a robot to acquire the spatial and wrench dynamics required for contact-rich manipulation
245 from a single demonstration Ijspeert et al. (2013); Fabisch (2024).

246 3.1.1 Cartesian DMP

247 We follow the formulation in Fabisch (2024), which supports Cartesian trajectories in $SE(3)$ and treats
248 orientation using unit quaternions to ensure smooth interpolation on the unit sphere.

The position $\mathbf{x} \in \mathbb{R}^3$ evolves according to the standard second-order dynamical system:

$$\tau \dot{\mathbf{x}} = \mathbf{v}_x \quad (1)$$

$$\tau \dot{\mathbf{v}}_x = \alpha_x (\beta_x (\mathbf{g}_x - \mathbf{x}) - \mathbf{v}_x) + \mathbf{f}_x(s) \quad (2)$$

where \mathbf{g}_x is the goal position, τ is the temporal scaling factor, α_x and β_x are positive gains, and $\mathbf{f}_x(s)$ is the nonlinear forcing term learned from demonstration and represented using weighted sum of radial basis functions (RBFs). Whereas, a forcing term is given by:

$$\mathbf{f}(s) = \frac{\sum_{i=1}^N \psi_i(s) w_i}{\sum_{i=1}^N \psi_i(s)}, \quad \psi_i(s) = \exp(-h_i(s - c_i)^2), \quad (3)$$

249 where w_i are the learned weights, c_i are the centers, and h_i are widths of the basis functions. The phase
250 variable s evolves over time according to the canonical system:

$$\tau \dot{s} = -\alpha_s s, \quad (4)$$

251 with $\alpha_s > 0$.

For orientation, let $\mathbf{q}(t) \in \mathbb{H}$ denote the unit quaternion representation, and $\boldsymbol{\omega} \in \mathbb{R}^3$ the angular velocity. Using the full quaternion error formulation Ude et al. (2014); Fabisch (2024), we write

$$\tau \dot{\boldsymbol{\omega}} = \alpha_{\omega} \delta(\mathbf{q}, \mathbf{q}_d) - \beta_{\omega} \boldsymbol{\omega} - \alpha_{\omega} \delta(\mathbf{q}_0, \mathbf{q}_d) s + \alpha_{\omega} \mathbf{f}_q(s), \quad (5)$$

$$\tau \dot{\mathbf{q}} = \frac{1}{2} \boldsymbol{\omega} \otimes \mathbf{q}. \quad (6)$$

252 Here, \mathbf{q}_0 and \mathbf{q}_d denote respectively the initial and desired (goal) quaternions, and \mathbf{f}_q is the respective
 253 forcing term. The positive constants α_{ω} and β_{ω} control the convergence rate and damping, respectively.
 254 The logarithmic quaternion error is

$$\delta(\mathbf{q}_1, \mathbf{q}_2) = 2 \log(\mathbf{q}_2 \otimes \bar{\mathbf{q}}_1), \quad (7)$$

255 with \otimes the quaternion product, $\bar{\mathbf{q}}_1$ the conjugate of \mathbf{q}_1 , and $\log(\cdot)$ the $\mathbb{S}^3 \rightarrow \mathfrak{so}(3)$ logarithmic map. This
 256 formulation constrains the trajectory to the unit quaternion manifold, avoiding singularities and enabling
 257 consistent learning and reproduction.

258 3.1.2 Wrench DMP

To model contact interaction, a similar parallel DMP is used for the wrench vector $\boldsymbol{\lambda} \in \mathbb{R}^6$:

$$\tau \dot{\boldsymbol{\lambda}} = \mathbf{v}_{\lambda} \quad (8)$$

$$\tau \dot{\mathbf{v}}_{\lambda} = \alpha_{\lambda} (\beta_{\lambda} (\mathbf{g}_{\lambda} - \boldsymbol{\lambda}) - \mathbf{v}_{\lambda}) + \mathbf{f}_{\lambda}(s), \quad (9)$$

259 where \mathbf{g}_{λ} is the goal wrench, and the forcing term \mathbf{f}_{λ} is learned componentwise for all six dimensions, with
 260 all parameters similarly defined as in Cartesian DMP. This DMP shares the same phase variable s as the
 261 Cartesian DMP to ensure time alignment between motion and wrench evolution Gams et al. (2014).

262 3.1.3 Training of wDMPs

263 For each novel assembly, a human performs kinesthetic teaching of the robot in gravity-compensation
 264 mode to provide a single demonstration. The measured pose and wrench trajectories are used to fit the
 265 forcing terms \mathbf{f}_x , \mathbf{f}_q , \mathbf{f}_{λ} (eq. 3) via imitation learning using ridge regression Hoerl and Kennard (2000). The
 266 operator segments the task into two stages, the (i) alignment, and (ii) insertion stage. One wDMP is learned
 267 per stage. At execution time, the Cartesian DMP yields the desired EE pose trajectory, while the Wrench
 268 DMP simultaneously produces the expected interaction wrench.

269 3.2 Contact Exploration via Force-driven Jiggling Motions

270 Having an expert-learned wDMP is not enough to guarantee a successful reproduction of the learned
 271 assembly. In reality, small misalignments during reproduction could occur due to modeling inaccuracies or
 272 minor geometric mismatches. To overcome this issue an error-driven, force-based exploration policy is
 273 employed between the alignment and insertion phase. To facilitate local contact exploration, a time-varying
 274 force $\mathbf{F}_{c,exp}(t) \in \mathbb{R}^3$ is applied at the EE. This force consists of sinusoidal components along each Cartesian
 275 direction:

$$F_{c,exp}^{(i)}(t) = A_i \sin(2\pi \mathbf{f}_i t) \cdot (d_{nom,i} + \delta \cdot \text{sign}(d_{nom,i})), \quad (10)$$

$$i \in \{x, y, z\}$$

276 where A_i and f_i are the amplitude and frequency of the sinusoidal force in the i -th direction, and $0 < \delta \ll 1$
 277 is a small perturbation term added to prevent the exploration force from vanishing near the goal. The
 278 probing direction \mathbf{d}_{nom} is computed as the normalized vector:

$$\mathbf{d}_{\text{nom}} = \frac{\mathbf{p}_g - \mathbf{p}_c}{\|\mathbf{p}_g - \mathbf{p}_c\|}, \quad (11)$$

279 where \mathbf{p}_g denotes the estimated goal position (e.g., hole center), and \mathbf{p}_c is the current EE position. Selecting
 280 low-frequency for these forces allows exploration of the contact surface, while high-frequency refines an
 281 existing contact. These micro-jiggling motions generate interaction forces that assist in resolving edge
 282 contacts and overcoming tight clearances.

283 Naturally, an alignment detection step is required to check if the exploration strategy resolved the
 284 misalignment. For this purpose, a similar jingling force is applied, but only in the assembly direction $\hat{\mathbf{n}}_d$
 285 given by wDMP. This is:

$$\mathbf{F}_{c,ad}(t) = \left(F_{c,ad}^{\min} + \frac{F_{c,ad}^{\max} - F_{c,ad}^{\min}}{1 + \exp(\alpha_F(c_F - \sin_F))} \right) \hat{\mathbf{n}}_d, \quad (12)$$

$$\sin_F = \frac{\sin(2\pi f_F t) + 1}{2}$$

286 where $F_{c,ad}^{\min}$, $F_{c,ad}^{\max}$ are the bounds of the applied force, α_F , c_F are logistic parameters, and f_F is
 287 its frequency. These parameters, along with the ones in 10, must be tuned empirically according to
 288 the desired outcome for the given robotic system. The corresponding contact torques are computed as,
 289 $\boldsymbol{\tau}_c = \mathbf{r}_{EE} \times \mathbf{F}_c \in \mathbb{R}^3$, where \mathbf{r}_{EE} is the position of the EE obtained from forward kinematics.

290 The exploration process continues until a maximum number of attempts is reached or a successful
 291 alignment is detected by the ART-based contact classifier described next.

292 3.3 Contact Classification via Distributed Dual Vigilance Fuzzy ART (DDVFA)

293 A typical ART cycle is managed by an orienting subsystem that enforces a “good-enough” similarity test.
 294 The working principle involves the following steps: (i) optionally encode the raw input in \mathbf{F}_0 field (layer),
 295 and place the encoded vector in \mathbf{F}_1 , the feature representation field; (ii) forward \mathbf{F}_1 to \mathbf{F}_2 , the category
 296 (class) field, where nodes compete and the best candidate is selected; (iv) check vigilance via a resonance
 297 test, matching the input with the expectation of an active \mathbf{F}_2 node; (v) if the match is sufficient, update the
 298 selected node, otherwise commit a novel one to incorporate the input data into long-term memory created
 299 by weight vectors between nodes of the \mathbf{F}_2 and \mathbf{F}_1 fields.

The DDVFA variant of ART implements a nested approach, in which the classes are represented by nodes that are themselves FuzzyART NNs. In FuzzyART, each input feature vector $\boldsymbol{\chi} \in \mathbb{R}^n$, $0 \leq \chi_i \leq 1 \forall i$ is *complement coded* when passing through the \mathbf{F}_0 field, i.e., $\mathbf{I} = (\boldsymbol{\chi}, \mathbf{1} - \boldsymbol{\chi})^T$. This vector is then forwarded by the \mathbf{F}_1 field to the \mathbf{F}_2 field. An orienting subsystem subsequently starts a competition between the \mathbf{F}_2 nodes. For this purpose, each node j determines its activation T_j based on the input and its weight vector \mathbf{w}_j which serves as the memory:

$$T_j = \frac{\|\mathbf{I} \wedge \mathbf{w}_j\|_1}{\alpha_T + \|\mathbf{w}_j\|_1}, \quad (13)$$

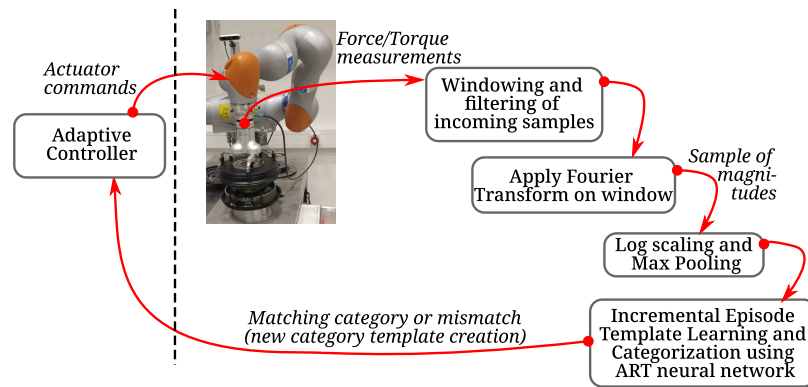


Figure 2. Dataflow of the ART-based contact learning and classification in the industrial assembly scenario

300 where $\|\cdot\|_1$ is the L_1 norm, \wedge denotes the fuzzy set AND operator (intersection, element-wise minimum),
 301 and α_T biases selection toward uncommitted nodes.

302 Nodes are considered in descending order of T_j until one satisfies the resonance criterion

$$M_j = \frac{\|\mathbf{I} \wedge \mathbf{w}_j\|_1}{\|\mathbf{I}\|_1} \geq \rho, \quad (14)$$

303 with vigilance parameter $\rho \in [0, 1]$ and match value M_j (similarity). The node j that passes the resonance
 304 test, updates its weights according to

$$\mathbf{w}_j^{\text{new}} = (1 - \eta) \mathbf{w}_j^{\text{old}} + \eta (\mathbf{I} \wedge \mathbf{w}_j^{\text{old}}), \quad (15)$$

305 with learning rate $\eta \in [0, 1]$. In DDVFA, a secondary, less strict global vigilance parameter extends this
 306 mechanism to regulate learning at the level of grouped classes.

307 In this work, DDVFA learns and classifies data from a stream of EE wrench samples composed of force
 308 and torque vectors, each $\in \mathbb{R}^3$. Before entering the ART network, the wrench stream is preprocessed as
 309 follows: (i) consecutive sample blocks are windowed with a Blackman–Harris window and transformed
 310 to the frequency domain via the Fast Fourier Transform (FFT); (ii) the number of features is reduced by
 311 max-pooling on the resulting frequency magnitudes; and (iii) all features are scaled to $[0, 1]$. The resulting
 312 feature vectors are then presented to the ART NN. The overall pipeline is illustrated in Fig. 2.

313 In our framework, ART is trained during the first assistive reproduction, after a wDMP has been learned.
 314 During the alignment phase, the human guides the robot slightly to align properly in case of millimeter errors.
 315 After correct alignment is ensured, 12 is applied to collect F/T data characterizing correct alignment with
 316 the assembly fixture. These data are used as input to the learning and classification procedure outlined above.
 317 This results in learning a set of template classes which map to episodes in the contacts of the assembly
 318 process. If the learning and adaption is further only activated during the successful alignment check, a
 319 distinct learning and further recognition of the frequency magnitude pattern of successful alignment can be
 320 achieved, mapping all other patterns to a mismatch (category index -1).

321 3.4 Adaptive Model Predictive Control (AMPC)

In this work, AMPC is used to reproduce the assembly task. The MPC problem is expressed in joint-space as in Rakovitis and Mronga (2024):

$$\begin{aligned}
 \min_{\mathbf{x}, \mathbf{u}} \quad & \|\mathbf{x}_d(T) - \mathbf{x}(T)\|_Q^2 + \int_0^T (\|\mathbf{x}_d - \mathbf{x}\|_Q^2 + \|\mathbf{u}\|_R^2) dt \\
 \text{s.t.} \quad & \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}) \\
 & \mathbf{x}(0) = \mathbf{x}_0 \\
 & \mathbf{x}_{\min} \leq \mathbf{x} \leq \mathbf{x}_{\max} \\
 & \mathbf{u}_{\min} \leq \mathbf{u} \leq \mathbf{u}_{\max}
 \end{aligned} \tag{16}$$

322 where, the system state is denoted by $\mathbf{x} = (\mathbf{q}, \dot{\mathbf{q}})$, encompassing joint positions and velocities; $\mathbf{u} = \boldsymbol{\tau}$
 323 represents the control input in the form of joint torques; \mathbf{x}_d denotes the joint-space reference obtained via
 324 Inverse Kinematics (IK) on the state-space reference given by the Cartesian DMP; \mathbf{x}_0 is the initial state;
 325 and T represents the prediction horizon. The constraints ensure that both the state and input remain within
 326 defined bounds \mathbf{x}_{\min} , \mathbf{x}_{\max} , \mathbf{u}_{\min} , \mathbf{u}_{\max} . Assuming a rigid contact at the EE, the dynamics \mathbf{f} are given by:

$$\mathbf{H}(\mathbf{q}) \ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) = \boldsymbol{\tau} + \mathbf{J}_c^\top \boldsymbol{\lambda}_c. \tag{17}$$

327 Here, $\mathbf{H}(\mathbf{q})$ is the inertia matrix in joint space, $\ddot{\mathbf{q}}$ is the joint acceleration, $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})$ represents the Coriolis,
 328 centrifugal, and gravitational effects, and \mathbf{J}_c is the Jacobian at the contact point. The contact wrench
 329 $\boldsymbol{\lambda}_c \in \mathbb{R}^6$ is modeled using a Cartesian impedance model (CIM):

$$\boldsymbol{\lambda}_c = \mathbf{K} \mathbf{e}_x + \mathbf{D} \dot{\mathbf{e}}_x + \boldsymbol{\lambda}_d, \tag{18}$$

330 where the stiffness and damping \mathbf{K} , $\mathbf{D} \in \mathbb{R}^{6 \times 6}$, respectively, are diagonal, positive semi-definite matrices.
 331 Those control the state-space compliance w.r.t. the EE pose and twist errors \mathbf{e}_x , $\dot{\mathbf{e}}_x$. The wrench $\boldsymbol{\lambda}_d$ is the
 332 desired wrench applied at the EE, provided in real-time either by the Wrench DMP or by the force-based
 333 contact exploration policy, according to the stage of the assembly. This formulation is chosen as it allows
 334 to control the compliance in state-space, while preserving postural joint-space control. Specifically, this
 335 compliance is adapted before solving MPC based on operation uncertainty.

336 3.5 Uncertainty Estimation

337 The uncertainty during execution of an assembly operation is derived using a GMM, trained during
 338 the human-assisted reproduction. After correct alignment is ensured by the operator, the insertion phase
 339 takes place. The observed Cartesian pose errors and wrenches recorded at the EE during this phase are
 340 collected into a dataset $\mathcal{D}^{(i)} = \{\mathbf{z}^{(i)} = [\mathbf{e}^{(i)}, \boldsymbol{\lambda}_m^{(i)}]\}_{i=1}^N$, where $\mathbf{e} = [\mathbf{e}_x^\top, \dot{\mathbf{e}}_x^\top]$ denotes the error measurements,
 341 and $\boldsymbol{\lambda}_m = [\mathbf{F}_m^\top, \boldsymbol{\tau}_m^\top]$ represents the F/T sensor measurements. This dataset is used to train a GMM
 342 with Expectation Maximization Dempster et al. (1977), to model the joint distribution of the error-F/T
 343 measurements. A GMM is defined as

$$p(\mathbf{z}) = \sum_{k=1}^K \pi_k \mathcal{N}_k(\mathbf{z} | \boldsymbol{\mu}_z^k, \boldsymbol{\Sigma}_z^k) \tag{19}$$

344 where \mathcal{N}_k denotes a multivariate Gaussian distribution with mean $\boldsymbol{\mu}_z^k$ and covariance matrix $\boldsymbol{\Sigma}_z^k$. Each
 345 component is weighted by a prior $\pi_k \in [0, 1]$, such that $\sum_{k=1}^K \pi_k = 1$. The number of mixture components,
 346 K , is treated as a hyperparameter and selected via grid search.

347 The objective is to use this GMM to derive an uncertainty metric, that describes the proximity to OOD
 348 measurements during the assembly process. The likelihood of an observation, given by

$$\mathcal{L} = \log(p(\mathbf{z})), \quad (20)$$

349 measures how likely the observation belongs to the GMM distribution. To obtain a normalized uncertainty
 350 value, the likelihood is mapped to a logistic score

$$\mathcal{S}(\mathbf{z}) = \frac{1}{1 + \exp(\beta(\mathcal{L}(\mathbf{z}) - c))}. \quad (21)$$

351 The center c and slope β are set by a calibration between in-distribution (ID) data of \mathcal{D} and a synthetic
 352 near-OOD set \mathcal{D}_{syn} , such that:

$$c = \frac{\tilde{\mathcal{L}}(\mathcal{D}) + \tilde{\mathcal{L}}(\mathcal{D}_{\text{syn}})}{2}, \quad \beta = \frac{\log\left(\frac{1}{\varepsilon} - 1\right)}{c - \tilde{\mathcal{L}}(\mathcal{D})}, \quad (22)$$

353 with $\varepsilon \rightarrow 0^+$. The set \mathcal{D}_{syn} is built by placing points on a χ^2 shell around each ID point (training samples),
 354 approximated using a responsibility-weighted local Gaussian Battistelli and Chisci (2014), so that synthetic
 355 samples lie in its negligible-probability tail. This yields low uncertainty on ID data and a smooth increase
 356 as the system drifts toward unmodeled regimes.

357 This uncertainty is used to modulate the MPC compliance in real-time, by adapting the stiffness and
 358 desired wrench of the CIM via logistic functions. The linear (l) and angular (a) stiffness of the CIM adapts,
 359 for each $i \in \{l, a\}$ as:

$$\mathbf{K}_i(\mathcal{S}) = \left(K_i^{\min} + \frac{K_i^{\max} - K_i^{\min}}{1 + \exp(\alpha_K(c_K - \mathcal{S}))} \right) \hat{\mathbf{n}}_{d,i}, \quad (23)$$

while the desired contact wrench adapts as

$$\boldsymbol{\lambda}_d(\mathcal{S}) = \boldsymbol{\lambda}_{\text{stage}} + \boldsymbol{\lambda}_{d,r}, \quad \boldsymbol{\lambda}_{d,r} = \begin{bmatrix} \mathbf{F}_{d,r} \\ \mathbf{r}_{EE} \times \mathbf{F}_{d,r} \end{bmatrix} \quad (24)$$

$$\mathbf{F}_{d,r} = \left(F_{d,r}^{\min} + \frac{F_{d,r}^{\max} - F_{d,r}^{\min}}{1 + \exp(\alpha_r(c_r - \mathcal{S}))} \right) \hat{\mathbf{n}}_d. \quad (25)$$

360 Here, $K_i^{\min}, K_i^{\max} \geq 0$ are the stiffness bounds, and the parameters $\alpha < 0, c \in [0, 1]$ denote the slope
 361 and midpoint of the logistics; $\boldsymbol{\lambda}_{\text{stage}}$ is the nominal wrench given by the Wrench DMP or the contact
 362 exploration, depending on the stage of the assembly; and $F_{d,r}^{\min}, F_{d,r}^{\max} \leq 0$ are the force bounds of the
 363 retraction wrench $\boldsymbol{\lambda}_{d,r}$. The desired bounds and sigmoid parameters are selected empirically to ensure
 364 stable and desired compliant behavior during contact. This formulation modulates compliance according
 365 to uncertainty, maximizing compliance when uncertainty is high (e.g. during unfamiliar contacts) and
 366 minimizing it when confidence is strong.




Figure 3. (Top) IndustRealKit parts, (Bottom) Disc brake parts

4 EXPERIMENTAL EVALUATION

367 The objective of this section is to evaluate the proposed approach on flexible assembly scenarios, where
 368 the assembled parts, while also their starting and goal configuration might constantly change. The target
 369 assembly tasks are defined as:

370 i) **IndustRealKit** (Tang et al., 2023), Fig. 3-top

- 371 • **3x Cylinder Pegs:** Small (8mm), Medium (12mm), Large (16mm) must be inserted into
 372 corresponding holes (clearances: 0.5-0.6mm),
- 373 • **3x Orthogonal Pegs:** Small (8mm), Medium (12mm), Large (16mm) must be inserted into
 374 corresponding holes (clearances: 0.5-0.6mm),
- 375 • **3x Gears:** Small (20mm), Medium (40mm), Large (60mm) must be inserted onto corresponding
 376 gearshafts (diametral clearances: 0.5mm),
- 377 • **2x Plug connectors:** i) 2-prong, and ii) 3-prong. Must be inserted into corresponding sockets,

378 ii) **Automotive (Car) parts (Disc Brake)**, Fig. 3-bottom

- 379 • **Wheel Bearing:** must be inserted onto the spindle,
- 380 • **Wheel Disc:** must be placed onto the wheel hub.

381 We train our framework on three parts only: (i) the large cylindrical peg, (ii) the 3-prong plug, and (iii) the
 382 wheel bearing. The goal is to generalize to related parts in each case respectively: (i) other peg-in-hole
 383 variants, (ii) the two-prong plug, and (iii) the wheel disc. We assume the operator selects the task class
 384 (peg-in-hole, plug-insertion, or disc brake assembly) in advance. For each case, the learned, task-specific
 385 wDMPs, ART, and GMM models are used.

386 For each learned assembly, we run multiple trials to reproduce all variations using each of the following
 387 controllers:

- 388 • **CIC:** a Cartesian Impedance Controller with joint limits and singularities avoidance utilizing the
 389 nullspace, as implemented in (Origanti et al., 2025). This controller maintains fixed empirical stiffness
 390 and does not use force-based contact exploration or alignment detection.




Figure 4. Output of force applied during contact exploration (eq. 10) with: $f_x = 2.7\text{Hz}$, $f_y = 5.4\text{Hz}$, $f_z = 4.5\text{Hz}$

- 391 • **MPC**: standard non-uncertain MPC. This method uses $K, D = 0$ in 18, and does not use force-based
 392 contact exploration or alignment detection.
- 393 • **MPVIC**: MPC with a Cartesian impedance model, which adapts stiffness based on pose errors (inspired
 394 from (Anand et al., 2023)) instead of uncertainty (increases stiffness while error gets larger). This
 395 methods sets $\lambda_d = \lambda_{\text{stage}}$ and skips the alignment detection phase.
- 396 • **uMPC-ART**: our approach, but without executing a retrieval on misalignment.
- 397 • **uMPC-ART_{r+}**: our approach with retrieval.
- 398 • **uMPC-ART_{r+}^{stiff}**: our approach with retrieval and with fixed maximum stiffness $K = K_{\text{max}}$, $D = D_{\text{max}}$,
 399 instead of adaptive via uncertainty.

400 For all the above, the MPC is solved at 200Hz using the Feasibility-driven Differential Dynamic
 401 Programming approach of the Crocoddyl library (Mastalli et al., 2020), which computes robot dynamics
 402 via the pinocchio library (Carpentier et al., 2019). The IK solution, computed with the TRAC-IK library
 403 (Beeson and Ames, 2015), is configured to find the closest solution to the current state and provides
 404 references to the MPC at 200Hz.

405 For contact exploration, we tune the frequency ratios in 10 to symmetrically excite the EE about the
 406 estimated goal in the XY plane—defined as the assembly (mating) surface—with Z aligned to the insertion
 407 axis (surface normal) (Fig. 4).

408 The ART classifier is parameterized by setting the number of samples in a window for the FFT, the
 409 overlapping, the size of the blocks used for max pooling, the local and global vigilance parameters, and the
 410 scaling to use. When started, it initializes by default with deactivated learning, allowing to pre-train a linear
 411 MinMaxScaler of the scikit-learn library (Pedregosa et al., 2011), when the encoded data's minimum and
 412 maximum cannot be estimated in advance. To obtain reasonable values for these parameters, a grid search
 413 can be performed on offline data. An exemplary result from such grid search based on a successful and a
 414 failed attempt to assemble the wheel bearing is shown in Fig. 5. In particular, learning is only activated
 415 on the first part of the data (until $t = 200\text{s}$) corresponding to the correct alignment, and tested with the
 416 remaining part of the data corresponding to a failed assembly. We observe that in the successful case,
 417 the part is initially jammed ($t \approx 140 - 155\text{s}$) causing high forces in Z-Axis, but the contact exploration
 418 resolves the misalignment at about $t \approx 155\text{s}$. During the alignment check that follows ($t \approx 170 - 185\text{s}$),
 419 the classifier learns the corresponding frequency magnitude pattern during correct alignment and maps it to
 420 a unique category 4. On the other hand, in the failed attempt to align the parts ($t \approx 265 - 275\text{s}$), neither
 421 this pattern nor the pattern of the prior exploration phase ($t \approx 245 - 255\text{s}$) is stably recognized due to the

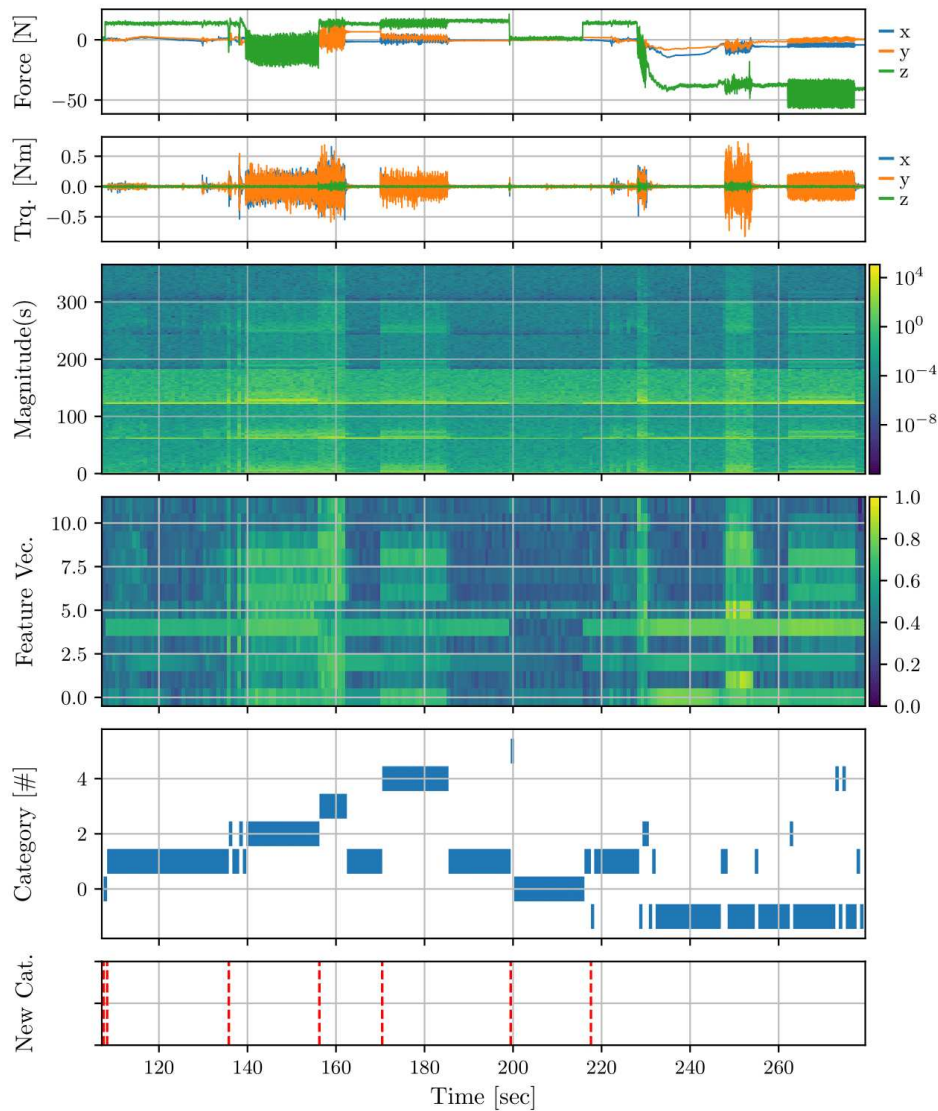


Figure 5. Exemplary data from the ART-based contact classification. From top to bottom: the six channels of the raw measurement (three forces, three torques), the corresponding frequency magnitudes stacked, the magnitudes after max-pooling and scaling, the assigned category index, and indication of novelty (mismatch).

422 dissimilarity to the pattern of correct alignment. This results in a significant number of mismatches, where
 423 the current input pattern does not match any of the existing categories. In such cases the classifier returns
 424 an identifier of -1 . Further, we use an additional median filter on the category identifier that is output
 425 by the classifier to handle spurious occurrences of patterns and instable class assignments. Equivalently,
 426 to only capture relevant patterns, the classifier can be operated such that learning is only active during
 427 an assembly or alignment check where correct alignment is ensured. In this mode of operation, all other
 428 patterns are thus mismatching, which simplifies further processing of the classifiers' output, as a dedicated
 429 labeling of the category identifiers is avoided.

430 4.1 Results

431 Tables 1 and 2 summarize performance across the three task families (peg-in-hole, plug-insertion, and
 432 disc brake assembly) under varying start and goal configurations. We report (i) the success rate, and (ii) the




Figure 6. Snapshots of three assembly tasks following our policy. (Top) Medium cylinder peg, (Middle) 3-prong plug, (Bottom) Wheel bearing.

433 average completion time computed over successful trials. The evaluation comprises three trials for each
 434 peg-in-hole instance (three for each cylindrical peg, three for each orthogonal, and three for each gear), five
 435 trials per plug type (2- and 3-prong), and five trials per car part (wheel bearing and wheel disc). We discuss
 436 the main trends evident in Tables 1, 2 below.

437 Overall, the proposed method markedly outperforms the baselines, exhibiting a clear robustness–speed
 438 trade-off. The CIC, MPC and MPVIC baselines rarely succeed on peg-in-hole (CIC: 3/27; MPC: 1/27;
 439 MPVIC: 2/27), while on plug and car-parts performance is mixed (respectively, CIC: 6/10 and 5/10; MPC:
 440 1/10 and 7/10; MPVIC: 5/10 and 9/10). In contrast, uMPC-ART generalizes well on all tasks, with 18/27
 441 on pegs, 5/10 on plugs, and 7/10 on car parts. Adding the retrial mechanism, uMPC-ART_{r+} yields the best
 442 reliability overall, achieving 21/27 on pegs, 8/10 on plugs, and 10/10 on car parts. This indicates that
 443 the ART-based misalignment detection followed by retrial, significantly improves success in uncertain
 444 tight-tolerance insertions tasks. The stiff variant, uMPC-ART_{r+}^{stiff}, further improves plugs to 9/10 but
 445 regresses on pegs to 18/27 and on car-parts to 8/10, suggesting that uncertainty-adaptive compliance is a
 446 better fit for diverse assembly tasks.

447 In terms of timing, when the baselines do succeed they are comparatively fast. CIC completes successful
 448 trials the quickest on average $\sim 30-46$ s, followed by MPC with $\sim 28-53$ s. MPVIC's few successful
 449 peg-in-hole trials complete around 53s, and on plugs/car parts it averages to $\sim 65-71$ s, which is comparable
 450 to uMPC-ART ($\sim 61-95$ s) but without the same reliability. Adding retrials, uMPC-ART_{r+} increases these
 451 to $\sim 74-118$ s (with Wheel Disc peaking at 139 ± 44.8 s); these longer horizons arise from retrials and
 452 the additional contact-exploration steps introduced by our policy. Nevertheless, uMPC-ART_{r+} improves
 453 timing relative to uMPC-ART_{r+}^{stiff} (\sim pegs: 87s, plugs: 126s, car-parts: 119s), indicating that adaptive
 454 compliance often benefits not only success but also completion time. This is a result of the fact that the
 455 stiff approach typically applies higher interaction forces, as seen in Fig. 7, which hinders misalignment
 456 resolution during contact exploration and thus leading more often to a retrial. In contrast, with adaptive
 457 compliance the interaction forces remain among the lowest in most trials and tasks, reducing them by
 458 $\sim 31\%$ vs. CIC, $\sim 53\%$ vs. MPC, $\sim 55\%$ vs. MPVIC, and $\sim 18\%$ vs. its stiff ablation, thereby enabling
 459 more efficient search of the contacted surface (Fig. 6) and, in turn, more frequent success.

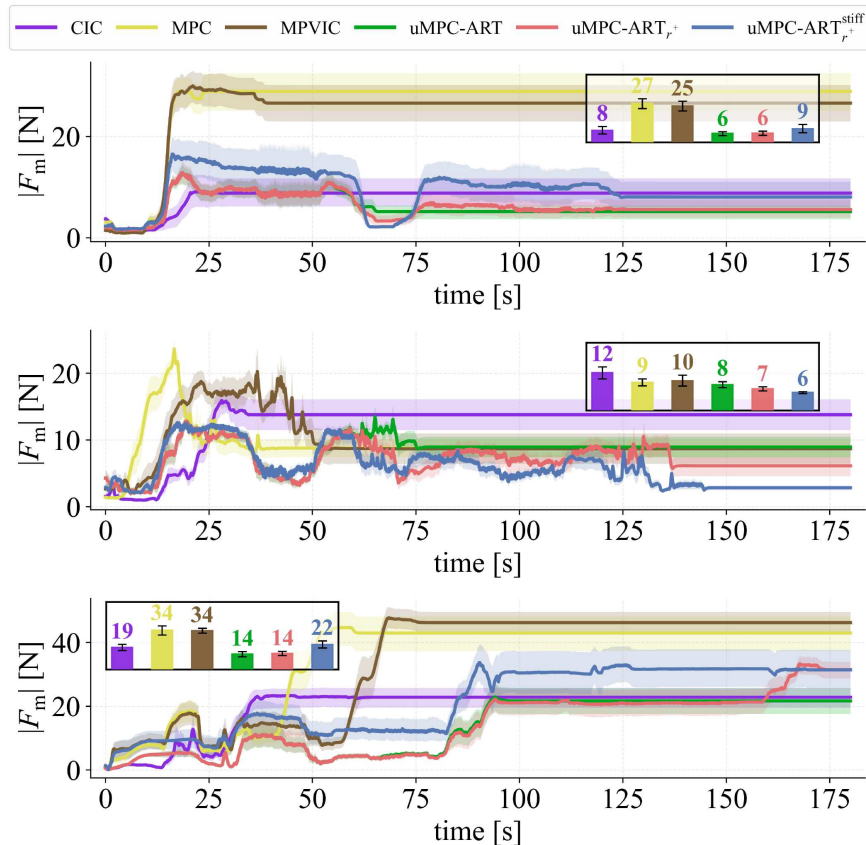


Figure 7. Force measurements over the 47 assembly trials listed in Tables 1, 2. **(Top)** 27x peg-in-hole, **(Middle)** 10x plug insertion, **(Bottom)** 10x car-parts. **(Inset)** Average force for all trials.

460 This trend can also be seen in Fig. 8, which plots the distance to the final goal pose over time for all
 461 tasks and trials. Independently of the parts, or start and goal configurations, the uMPC-ART_{r+} variant (red)
 462 maintains the lowest steady-state residuals with tight variability. Although the retrieval policy produces brief
 463 mid-trajectory spikes, corresponding to the automatic retreat and re-attempt, these are followed by renewed
 464 convergence. Notably $\text{uMPC-ART}_{r+}^{\text{stiff}}$ (blue) triggers retrievals more often, as evidenced by larger retreat
 465 curves and greater dispersion than uMPC-ART_{r+} .

466 Aggregating across all tasks in Table 3, the trend is unambiguous: uMPC-ART_{r+} achieves the highest
 467 overall success (83.0%), followed by the stiff variant, $\text{uMPC-ART}_{r+}^{\text{stiff}}$ (74.5%), and non-retrieval uMPC-ART
 468 (63.8%), while MPVIC (34.0%), CIC (29.8%), and MPC (19.1%) trail. Taken together, the results support
 469 three takeaways: (i) uncertainty-based adaptive compliance is key for generalization across shapes and
 470 pose variations; (ii) a retrieval policy after misalignment detection substantially boosts reliability at the cost
 471 of longer executions; and (iii) for deployments prioritizing success under variability, uMPC-ART_{r+} is the
 472 most reliable default, but when cycle time is paramount and parts are less ambiguous, uMPC-ART without
 473 retrieval offers a favorable speed–reliability balance.

5 CONCLUSION & OUTLOOK

474 To summarize, we presented a unified, data-efficient framework for fast programming and adaptive
 475 reproduction of contact-rich assembly tasks that combines synchronized wrench–motion DMPs with
 476 a GMM-based, uncertainty-aware AMPC, and an ART contact classifier. Training requires only two

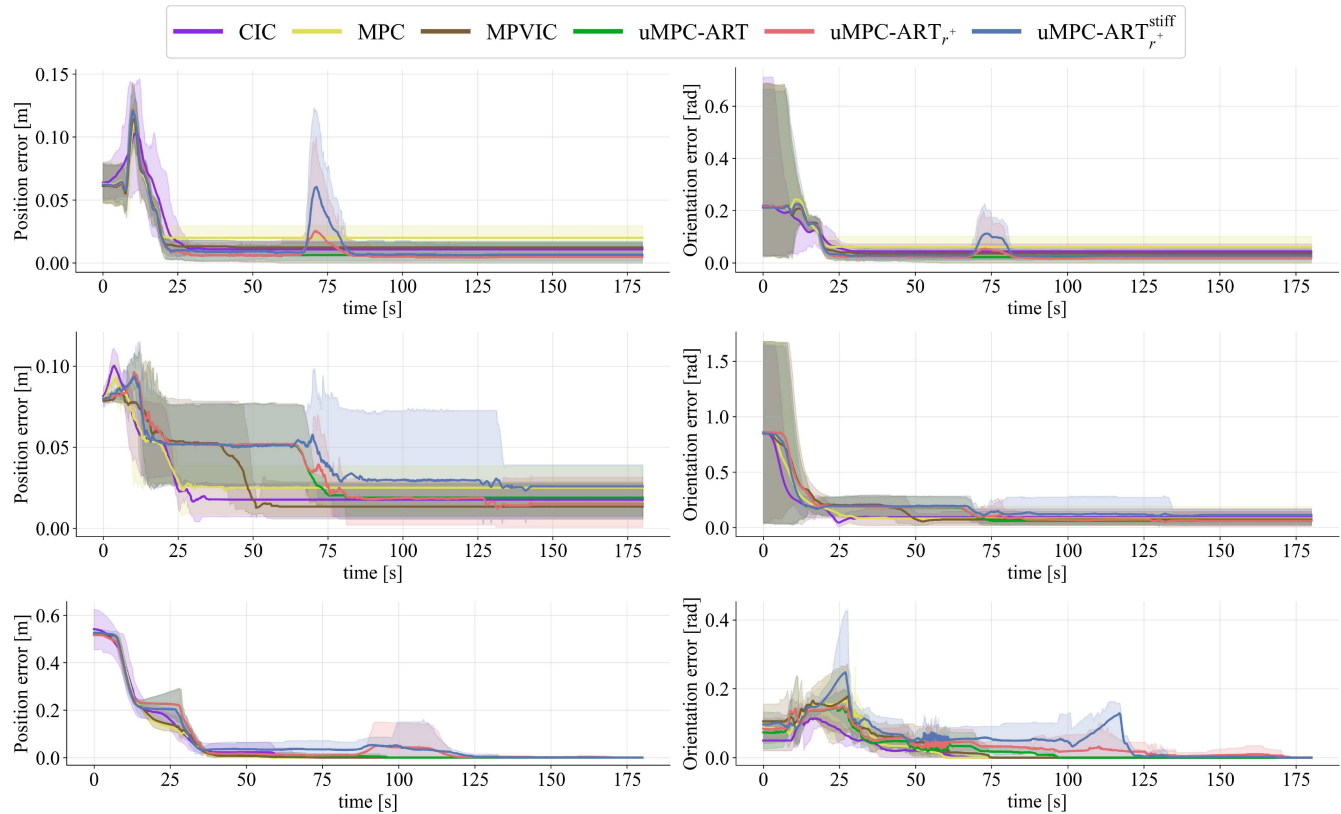


Figure 8. Distance to final position (Left) and orientation (Right) goals. Results over the 47 assembly trials listed in Tables 1, 2: **(Top)** 27x peg-in-hole, **(Middle)** 10x plug insertion, **(Bottom)** 10x car-parts.

477 demonstrations: (i) kinesthetic teaching of wrench and motion profiles, and (ii) an assistive reproduction in
 478 which the operator ensures alignment so the robot can collect contact data for ART (alignment detection)
 479 and the GMM (uncertainty estimation). At run time, the robot adapts compliance via the GMM likelihood,
 480 which is key for enabling adaptive exploration of the contact surface using force-based jiggling motions,
 481 and schedules assembly stage transitions (e.g. retrieval on misalignment) based on the contact context inferred
 482 by ART.

483 Trained on a single instance each of peg-in-hole, plug insertion, and wheel-bearing assembly, the system
 484 successfully completed 39/47 ($\approx 83\%$) previously unseen assemblies with varied geometries and start/goal
 485 poses. It consistently outperformed classic CIC/MPC and an error-driven MPVIC baseline in success rate
 486 (reliability) at the cost of slightly increased completion time, while maintaining the lowest interaction forces
 487 overall. These results make a step towards flexible manufacturing lines where production requirements
 488 can change frequently. Although, we demonstrate the feasibility of such solution using only a F/T sensor
 489 and proprioception, future work should focus on advancing the autonomy of the system towards online
 490 learning, fusing vision or tactile information, extending to multi-contact or multi-stage assemblies with
 491 regrasping, and eventually minimizing manual effort.

FUNDING

492 This research was done in the HARTU project (funded by the European Union, grant number 101092100).

CONFLICT OF INTEREST

493 The authors declare that the research was conducted in the absence of any commercial or financial
494 relationships that could be construed as a potential conflict of interest.

GENERATIVE AI STATEMENT

495 The author(s) declare that no Generative AI was used in the creation of this manuscript.

DATA AVAILABILITY STATEMENT

496 The original contributions presented in the study are included in the article/supplementary material, further
497 inquiries can be directed to the corresponding author.

REFERENCES

- 498 Anand, A. S., Gravdahl, J. T., and Abu-Dakka, F. J. (2023). Model-based variable impedance learning
499 control for robotic manipulation. *Robotics and Autonomous Systems* 170, 104531
- 500 Anne, T., Wilkinson, J., and Li, Z. (2021). Meta-learning for fast adaptive locomotion with uncertainties in
501 environments and robot dynamics. In *2021 IEEE/RSJ International Conference on Intelligent Robots
502 and Systems (IROS)*. 4568–4575
- 503 Arcari, E., Minniti, M. V., Scampicchio, A., Carron, A., Farshidian, F., Hutter, M., et al. (2023). Bayesian
504 multi-task learning mpc for robotic mobile manipulation. *IEEE Robotics and Automation Letters* 8,
505 3222–3229
- 506 Bargsten, V. and Kirchner, F. (2023). Actuator-level motion and contact episode learning and classification
507 using adaptive resonance theory. *Intelligent Service Robotics* doi:10.1007/s11370-023-00481-7
- 508 Bargsten, V., Rakovitis, D., Origanti, V. K., Danzglock, A., and Kirchner, F. (2025). Continuous
509 Learning of Contact Episodes from Proprioceptive Sensors in Industrial Assembly Scenarios Using
510 Adaptive Resonance Theory. In *Intelligent and Fuzzy Systems*, eds. C. Kahraman, S. Cebi, B. Oztaysi,
511 S. Cevik Onar, C. Tolga, I. Ucal Sari, and İ. Otay (Cham: Springer Nature Switzerland), 342–349.
512 doi:10.1007/978-3-031-98304-7_39
- 513 Battistelli, G. and Chisci, L. (2014). Kullback–leibler average, consensus on probability densities, and
514 distributed state estimation with guaranteed stability. *Automatica* 50, 707–718
- 515 Bednarczyk, M., Omran, H., and Bayle, B. (2020). Model predictive impedance control. In *2020 IEEE
516 International Conference on Robotics and Automation (ICRA)*. 4702–4708
- 517 Beeson, P. and Ames, B. (2015). Trac-ik: An open-source library for improved solving of generic inverse
518 kinematics. In *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*.
519 928–935
- 520 Brito da Silva, L. E., Elnabarawy, I., and Wunsch, D. C. (2020). Distributed dual vigilance fuzzy adaptive
521 resonance theory learns online, retrieves arbitrarily-shaped clusters, and mitigates order dependence.
522 *Neural Networks* 121, 208–228. doi:10.1016/j.neunet.2019.08.033
- 523 Calinon, S. (2016). A tutorial on task-parameterized movement learning and retrieval. *Intelligent Service
524 Robotics* 9, 1–29
- 525 Calinon, S., Guenter, F., and Billard, A. (2007). On learning, representing, and generalizing a task in
526 a humanoid robot. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 37,
527 286–298

- 528 Carpenter, G., Grossberg, S., and Rosen, D. (1991). Fuzzy ART: An adaptive resonance algorithm for
529 rapid, stable classification of analog patterns. In *IJCNN-91-Seattle International Joint Conference on*
530 *Neural Networks* (Seattle, WA, USA: IEEE), vol. ii, 411–416. doi:10.1109/IJCNN.1991.155368
- 531 Carpenter, G. A. and Grossberg, S. (1987). A massively parallel architecture for a self-organizing
532 neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing* 37, 54–115.
533 doi:10.1016/S0734-189X(87)80014-2
- 534 Carpentier, J., Saurel, G., Buondonno, G., Mirabel, J., Lamiroux, F., Stasse, O., et al. (2019). The pinocchio
535 c++ library – a fast and flexible implementation of rigid body dynamics algorithms and their analytical
536 derivatives. In *IEEE International Symposium on System Integrations (SII)*
- 537 Chang, C., Haninger, K., Shi, Y., Yuan, C., Chen, Z., and Zhang, J. (2022). Impedance adaptation by
538 reinforcement learning with contact dynamic movement primitives. In *2022 IEEE/ASME International*
539 *Conference on Advanced Intelligent Mechatronics (AIM)*. 1185–1191
- 540 Cho, C.-N., Kim, J.-H., Kim, Y.-L., Song, J.-B., and Kyung, J.-H. (2012). Collision Detection Algorithm
541 to Distinguish Between Intended Contact and Unexpected Collision. *Advanced Robotics* 26, 1825–1840.
542 doi:10.1080/01691864.2012.685259
- 543 Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the
544 em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* 39, 1–38
- 545 Fabisch, A. (2024). movement_primitives: Imitation learning of cartesian motion with movement primitives.
546 *Journal of Open Source Software* 9, 6695
- 547 Gams, A., Nemec, B., Ijspeert, A., and Ude, A. (2014). Coupling movement primitives: Interaction with
548 the environment and bimanual tasks. *IEEE Transactions on Robotics* 30, 816–830
- 549 Goyal, A., Blukis, V., Xu, J., Guo, Y., Chao, Y.-W., and Fox, D. (2024). RVT-2: Learning precise
550 manipulation from few demonstrations. In *Accepted at Robotics: Science and Systems (RSS) 2024*.
551 (forthcoming)
- 552 Grossberg, S. (1976). Adaptive pattern classification and universal recoding: II. Feedback, expectation,
553 olfaction, illusions. *Biological Cybernetics* 23, 187–202. doi:10.1007/BF00340335
- 554 Guo, Y., Tang, B., Akinola, I., Fox, D., Gupta, A., and Narang, Y. (2025). SRSA: Skill retrieval and
555 adaptation for robotic assembly tasks. In *Proceedings of the International Conference on Learning*
556 *Representations (ICLR)*. (forthcoming)
- 557 Hoerl, A. E. and Kennard, R. W. (2000). Ridge regression: Biased estimation for nonorthogonal problems.
558 *Technometrics* 42, 80–86
- 559 Hoffmann, H., Pastor, P., Park, D.-H., and Schaal, S. (2009). Biologically-inspired dynamical systems
560 for movement generation: Automatic real-time goal adaptation and obstacle avoidance. In *2009 IEEE*
561 *international conference on robotics and automation (IEEE)*, 2587–2592
- 562 Ijspeert, A. J., Nakanishi, J., Hoffmann, H., Pastor, P., and Schaal, S. (2013). Dynamical movement
563 primitives: Learning attractor models for motor behaviors. *Neural Computation* 25, 328–373
- 564 Iskandar, M., Albu-Schäffer, A., and Dietrich, A. (2024). Intrinsic sense of touch for intuitive physical
565 human-robot interaction. *Science Robotics* 9, eadn4008. doi:10.1126/scirobotics.adn4008
- 566 Iwata, T. and Kumagai, A. (2022). Meta-learning for out-of-distribution detection via density estimation in
567 latent space. *CoRR* abs/2206.09543
- 568 Jha, D. K., Romeres, D., Yezazunis, W., and Nikovski, D. (2022). Imitation and supervised learning of
569 compliance for robotic assembly. In *2022 European Control Conference (ECC)*. 1882–1889
- 570 Kim, Y.-L., Ahn, K.-H., and Song, J.-B. (2020). Reinforcement learning based on movement primitives for
571 contact tasks. *Robotics and Computer-Integrated Manufacturing* 62, 101863

- 572 Kouris, A., Dimeas, F., and Aspragathos, N. (2016). Contact distinction in human-robot cooperation with
573 admittance control. In *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*
574 (Budapest, Hungary: IEEE), 001951–001956. doi:10.1109/SMC.2016.7844525
- 575 Kouris, A., Dimeas, F., and Aspragathos, N. (2018). A frequency domain approach for contact type
576 distinction in human–robot collaboration. *IEEE Robotics and Automation Letters* 3, 720–727. doi:10.
577 1109/LRA.2017.2789249
- 578 Kramberger, A., Gams, A., Nemeč, B., Chrysostomou, D., Madsen, O., and Ude, A. (2017). Generalization
579 of orientation trajectories and force-torque profiles for robotic assembly. *Robotics and Autonomous*
580 *Systems* 98, 333–346
- 581 Leonardo Enzo Brito da Silva, Elnabarawy, I., and Wunsch II, D. C. (2019). A Survey of Adaptive
582 Resonance Theory Neural Network Models for Engineering Applications. *arXiv:1905.11437 [cs, stat]*
583 doi:10.48550/arXiv.1905.11437
- 584 Levine, S., Wagener, N., and Abbeel, P. (2015). Learning contact-rich manipulation skills with guided
585 policy search. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*. 156–163
- 586 Liu, J., Chang, Q., Xiao, G., and Biller, S. (2012). The costs of downtime incidents in serial multistage
587 manufacturing systems. *Journal of Manufacturing Science and Engineering* 134, 021016
588 [Dataset] Liu, J. and Chen, B. (2024). SonicSense: Object Perception from In-Hand Acoustic Vibration.
589 doi:10.48550/arXiv.2406.17932
- 590 Luo, J., Solowjow, E., Wen, C., Ojea, J. A., Agogino, A. M., Tamar, A., et al. (2019). Reinforcement
591 learning on variable impedance controller for high-precision robotic assembly. In *2019 International*
592 *Conference on Robotics and Automation (ICRA)*. 3080–3087
- 593 Mastalli, C., Budhiraja, R., Merkt, W., Saurel, G., Hammoud, B., Naveau, M., et al. (2020). Crocodyl: An
594 Efficient and Versatile Framework for Multi-Contact Optimal Control. In *IEEE International Conference*
595 *on Robotics and Automation (ICRA)*
- 596 Minniti, M. V., Grandia, R., Fäh, K., Farshidian, F., and Hutter, M. (2021). Model predictive robot-
597 environment interaction control for mobile manipulation tasks. In *2021 IEEE International Conference*
598 *on Robotics and Automation (ICRA)*. 1651–1657
- 599 Morgan, A. S., Bateux, Q., Hao, M., and Dollar, A. M. (2023). Towards generalized robot assembly
600 through compliance-enabled contact formations. In *2023 IEEE International Conference on Robotics*
601 *and Automation (ICRA)*. 8010–8016
- 602 Noseworthy, M., Tang, B., Wen, B., Handa, A., Kessens, C., Roy, N., et al. (2025). Forge: Force-guided
603 exploration for robust contact-rich manipulation under uncertainty. *IEEE Robotics and Automation*
604 *Letters* 10, 4436–4443
- 605 Origanti, V. K., Danzglock, A., and Kirchner, F. (2025). Look ahead optimization for managing nullspace
606 in cartesian impedance control of dual-arm robots. In *2025 IEEE/SICE International Symposium on*
607 *System Integration (SII)*. 990–997
- 608 Origanti, V. K., Eiband, T., and Lee, D. (2022). Automatic parameterization of motion and force controlled
609 robot skills. In *Robot Intelligence Technology and Applications 6* (Springer International Publishing),
610 66–78
- 611 Pankert, J. and Hutter, M. (2023). Learning Contact-Based State Estimation for Assembly Tasks. In *2023*
612 *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Detroit, MI, USA: IEEE),
613 5087–5094. doi:10.1109/IROS55552.2023.10342219
- 614 Paraschos, A., Daniel, C., Peters, J. R., and Neumann, G. (2013). Probabilistic movement primitives.
615 *Advances in Neural Information Processing Systems* 26

- 616 Park, H., Park, J., Lee, D.-H., Park, J.-H., Baeg, M.-H., and Bae, J.-H. (2017). Compliance-based robotic
617 peg-in-hole assembly strategy without force feedback. *IEEE Transactions on Industrial Electronics* 64,
618 6299–6309
- 619 Pastor, P., Hoffmann, H., Asfour, T., and Schaal, S. (2009). Learning and generalization of motor skills
620 by learning from demonstration. In *2009 IEEE international conference on robotics and automation*
621 (IEEE), 763–768
- 622 Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn:
623 Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830
- 624 Rakovitis, D. and Mronga, D. (2024). Gaussian mixture likelihood-based adaptive mpc for interactive
625 mobile manipulators. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*.
626 1392–1398
- 627 Saviolo, A., Frey, J., Rathod, A., Diehl, M., and Loianno, G. (2024). Active learning of discrete-time
628 dynamics for uncertainty-aware model predictive control. *IEEE Transactions on Robotics* 40, 1273–1291
- 629 Schoettler, G., Nair, A., Ojea, J. A., Levine, S., and Solowjow, E. (2020). Meta-reinforcement learning for
630 robotic industrial insertion tasks. In *2020 IEEE/RSJ International Conference on Intelligent Robots and*
631 *Systems (IROS)*. 9728–9735
- 632 Tang, B., Lin, M. A., Akinola, I., Handa, A., Sukhatme, G. S., Ramos, F., et al. (2023). Industreal:
633 Transferring contact-rich assembly tasks from simulation to reality. In *Robotics: Science and Systems*
- 634 Thelenberg, N. and Ott, C. (2024). On handling variable stiffness parameters in compliance control via
635 mpc. In *2024 European Control Conference (ECC)*. 615–620
- 636 Ude, A., Gams, A., Asfour, T., and Morimoto, J. (2010). Task-specific generalization of discrete and
637 periodic dynamic movement primitives. *IEEE Transactions on Robotics* 26, 800–815
- 638 Ude, A., Nemeč, B., Petrič, T., and Morimoto, J. (2014). Orientation in cartesian space dynamic movement
639 primitives. In *2014 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE),
640 2997–3004
- 641 Wu, Y., Chen, Z., Wu, F., Chen, L., Zhang, L., Bing, Z., et al. (2025). Tacdiffusion: Force-domain diffusion
642 policy for precise tactile manipulation. In *Proceedings of the IEEE International Conference on Robotics*
643 *and Automation (ICRA)*. (forthcoming)
- 644 Xu, J., Zheng, Y., Jiang, X., Yang, S., Xiang, L., and Zhang, Z. (2022a). Real-time inertial parameter
645 identification of floating-base robots through iterative primitive shape division. In *2022 International*
646 *Conference on Robotics and Automation (ICRA)*. 5960–5966
- 647 Xu, S., Zhu, L., and Ho, C. P. (2022b). Learning efficient and robust multi-modal quadruped locomotion:
648 A hierarchical approach. In *2022 International Conference on Robotics and Automation (ICRA)*.
649 4649–4655
- 650 Yan, C., Wu, J., and Zhu, Q. (2021). Learning-based contact status recognition for peg-in-hole assembly.
651 In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 6003–6009
- 652 Zong, B., Song, Q., Min, M. R., Cheng, W., Lumezanu, C., ki Cho, D., et al. (2018). Deep autoencoding
653 gaussian mixture model for unsupervised anomaly detection. In *International Conference on Learning*
654 *Representations*

Table 1. Peg-in-hole. Each cell shows **success rate (top)** and **completion time [s] of successful trials (bottom)**. Subtotal rows aggregate across all nine tasks. The superscript $\#r$ denotes the number of retrievals performed.

Method	Cylinder			Orthogonal			Gears		
	L	M	S	L	M	S	L	M	S
CIC	0/3	0/3	0/3	0/3	0/3	0/3	0/3	1/3	2/3
<i>Subtotal: (3/27, 30.3 ± 1.2s)</i>								31.0 ± 0.0s	30.0 ± 1.4s
MPC	0/3	0/3	0/3	0/3	0/3	0/3	0/3	1/3	0/3
<i>Subtotal: (1/27, 28.0 ± 0.0s)</i>								28.0 ± 0.0s	
MPVIC	0/3	0/3	0/3	0/3	0/3	0/3	0/3	0/3	2/3
<i>Subtotal: (2/27, 52.5 ± 2.1s)</i>									52.5 ± 2.1s
uMPC-ART	3/3	2/3	2/3	3/3	2/3	0/3	3/3	0/3	3/3
<i>Subtotal: (18/27, 61.1 ± 2.5s)</i>	57.3 ± 1.2s	61.5 ± 2.1s	62.0 ± 1.4s	64.0 ± 1.7s	62.0 ± 0.0s	-	59.3 ± 1.2s	-	62.0 ± 1.0s
uMPC-ART_{r+}	3/3 ^{#r=0}	3/3 ^{#r=1}	3/3 ^{#r=1}	3/3 ^{#r=0}	3/3 ^{#r=1}	0/3 ^{#r=3}	3/3 ^{#r=1}	0/3 ^{#r=3}	3/3 ^{#r=0}
<i>Subtotal: (21/27^{#r=10}, 73.9 ± 27.2s)</i>	57.3 ± 1.2s	81.0 ± 33.8s	84.3 ± 38.7s	64.0 ± 1.7s	89.0 ± 46.8s	-	79.7 ± 34.1s	-	62.0 ± 1.0s
uMPC-ART_{r+}^{stiff}	2/3 ^{#r=2}	3/3 ^{#r=0}	1/3 ^{#r=2}	3/3 ^{#r=0}	3/3 ^{#r=1}	1/3 ^{#r=3}	3/3 ^{#r=2}	0/3 ^{#r=3}	2/3 ^{#r=1}
<i>Subtotal: (18/27^{#r=14}, 87.3 ± 30.4s)</i>	99.0 ± 41.0s	65.0 ± 1.7s	128.0 ± 0.0s	63.3 ± 2.9s	104.7 ± 35.3s	119.0 ± 0.0s	101.3 ± 35.9s	-	62.0 ± 2.8s

Table 2. Plug insertion and car-part assembly. Each cell shows **success rate (top)** and **completion time [s] of successful trials (bottom)**. Subtotal rows aggregate the two Plug tasks and the two Car-parts tasks respectively. The superscript $\#r$ denotes the number of retrievals performed.

Method	Plug		Car parts	
	3-prong	2-prong	Wheel bearing	Wheel disc
CIC	4/5	2/5	2/5	3/5
<i>Subtotal</i>	33.8 ± 1.0s	36.5 ± 2.1s	55.5 ± 4.9s	40.3 ± 8.5s
<i>Subtotal</i>	(6/10, 34.7 ± 1.9s)		(5/10, 46.4 ± 10.5s)	
MPC	1/5	0/5	2/5	5/5
<i>Subtotal</i>	30.0 ± 0.0s	-	59.5 ± 0.7s	51.0 ± 0.7s
<i>Subtotal</i>	(1/10, 30.0 ± 0.0s)		(7/10, 53.4 ± 4.2s)	
MPVIC	4/5	1/5	4/5	5/5
<i>Subtotal</i>	65.5 ± 3.3s	64.0 ± 0.0s	74.3 ± 2.5s	67.6 ± 0.5s
<i>Subtotal</i>	(5/10, 65.2 ± 2.9s)		(9/10, 70.6 ± 3.8s)	
uMPC-ART	5/5	0/5	5/5	2/5
<i>Subtotal</i>	78.6 ± 4.6s	-	96.4 ± 0.5s	90.0 ± 0.0s
<i>Subtotal</i>	(5/10, 78.6 ± 4.6s)		(7/10, 94.6 ± 3.2s)	
uMPC-ART_{r+}	5/5 ^{#r=0}	3/5 ^{#r=5}	5/5 ^{#r=0}	5/5 ^{#r=3}
<i>Subtotal</i>	78.6 ± 4.6s	151.7 ± 17.6s	96.4 ± 0.5s	139.0 ± 44.8s
<i>Subtotal</i>	(8/10 ^{#r=5} , 106.0 ± 39.1s)		(10/10 ^{#r=3} , 117.7 ± 37.3s)	
uMPC-ART_{r+}^{stiff}	4/5 ^{#r=3}	5/5 ^{#r=5}	3/5 ^{#r=1}	5/5 ^{#r=1}
<i>Subtotal</i>	106.8 ± 31.7s	141.0 ± 5.2s	108.3 ± 21.6s	124.8 ± 33.1s
<i>Subtotal</i>	(9/10 ^{#r=8} , 125.8 ± 26.7s)		(8/10 ^{#r=2} , 118.6 ± 28.8s)	

Table 3. Overall success across all tasks

Method	Overall success
CIC	14/47 (29.8%)
MPC	9/47 (19.1%)
MPVIC	16/47 (34.0%)
uMPC-ART	30/47 (63.8%)
uMPC-ART _{r+}	39/47 (83.0%)
uMPC-ART _{r+} ^{stiff}	35/47 (74.5%)