

Beyond the Teacher: Leveraging Mixed-Skill Demonstrations for Robust Imitation Learning

S. Saharsh*, Shubham Sonkar*, Pushpak Jagtap, and Ravi Prakash

Abstract—Achieving expert-like robotic task execution in dynamic environments typically requires extensive, high-quality expert demonstrations, a significant bottleneck for real-world deployment. We present a novel learning framework that overcomes this data dependency, enabling robots to perform complex periodic tasks with expert-like proficiency, even when learning from naive demonstrations. Our two-stage pipeline first selects a representative demonstration based on user-defined information-aware task intention scores. This single best demo is then used to extract a canonical motion shape via Periodic Dynamic Movement Primitives (DMPs). Finally, a Long Short-Term Memory (LSTM) network refines the entire set of demonstrations, leveraging a multi-objective score that combines the canonical shape with mutual information and other task quality metrics. The proposed approach is demonstrated on a Franka Research 3 robot performing phasic tasks across three contrasting domains: wiping in human assistive services, weaving in the textile industry, and pick-and-place operations for warehouse automation. Visit project page at: <https://focaslabs.github.io/beyondtheteacher/>.

I. INTRODUCTION

Imitation Learning (IL) enables robots to acquire complex behaviors by learning state-action mappings from expert demonstrations, bypassing the need for hand-crafted reward functions or extensive environmental exploration. However, a major limitation of conventional IL methods is their dependence on large quantities of high-quality expert demonstrations. When demonstrations are suboptimal, inconsistent, or scarce, as is often the case in real-world robotic settings, learned policies suffer from compounding errors due to covariate shift, and scalability becomes a bottleneck due to high data acquisition costs.

To address this, we propose a novel IL pipeline (Fig. 1) that bypasses reliance on abundant expert data. Instead, our method recovers a high-quality, expert-like dataset from a single clean demonstration by combining periodic Dynamic Movement Primitives (DMPs) with LSTM-based trajectory sequence correction. This recovered dataset enables robust policy learning even when the remaining data is noisy, suboptimal, or corrupted.

Recent works such as Making Imitation Learning Easy with Self-Supervision (MILES) [1], Co-Imitation Learning (CoIL) [2], and Task-Oriented Self-Imitation Learning (TOSIL) [3] mitigate the disadvantage of poor demonstrations by using self-generated data through good or bad

* The authors have contributed equally.

This work was supported in part by the ARTPARK, IISc Bengaluru.

All authors are with the Cyber Physical Systems, Indian Institute of Science, Bengaluru, 560012, India. (e-mail: {saharsh2021, shubhamsonka, pushpak, ravipr}@iisc.ac.in)

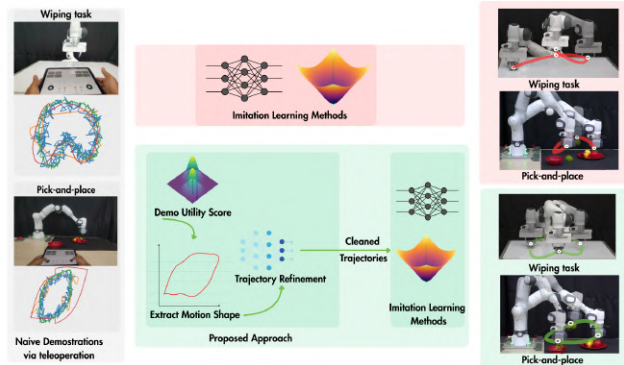


Fig. 1: Novel IL pipeline for Imitation learning with expert-like behavior using Imperfect demonstrations without human intervention. Results from standard IL methods are shown in red and from our proposed pipeline are shown in green.

experiences or human interventions. Others, like ILEED [4], estimate the ability of the demonstrator or trajectory feasibility to filter or weight demonstrations for behavior cloning. However, these methods typically require access to large demonstration datasets and sufficient exploration of the state space to effectively learn latent representations and struggle with complex, highly structured behaviors, such as periodic motions with high curvature, as seen in handwriting, wiping, or pick-and-place tasks.

Data curation methods for imitation learning under imperfect demonstrations, such as [9], introduce SCIZOR, a self-supervised framework designed for large-scale dataset refinement. SCIZOR filters out redundant and low-quality samples; however, it does not explicitly quantify the information content of state-action pairs, which may lead to the removal of task-relevant features. In [7], the authors analyze data quality by characterizing diversity and action consistency as key factors for generalization, yet lack an online mechanism to score samples. To address limited labeled data, [8] proposes retrieving behaviors from unlabeled datasets, a strategy that depends on sufficiently diverse prior data and effective retrieval metrics. Similarly, [6] demonstrates that incorporating fully observable policies can enhance learning in partially observable settings, emphasizing the importance of preserving predictive information. Authors in [10] employ mutual information estimators to measure state-action dependence, showing that greater mutual information is associated with improved imitation performance.

Taken together as in Table I, these works indicate two contrasting directions in data curation: first, where ap-

Challenge	P1[2]	P2[5]	P3[6]	P4[4]	P5[7]	P6[8]	P7[3]	P8[1]	P9[9]	P10[10]	P11
Robust to Imperfect Demonstrations	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Learns from Few Expert Trajectories	✗	✗	✓	✗	✗	✗	✗	✓	✗	✗	✓
Handles Periodic or Structured Motions	✗	✓	✓	✗	✓	✗	✗	✗	✗	✗	✓
Preserves Task-Relevant Information	✓	✗	✗	✓	✗	✗	✗	✓	✓	✓	✓
Multi-Objective Loss or MI-Based Filtering	✗	✗	✗	✓	✓	✗	✗	✗	✓	✓	✓

TABLE I: Comparison of Imitation Learning methods on key challenges in Imitation Learning from imperfect demonstrations. In this table, ✓ indicates that the method addresses the challenge, while ✗ indicates that the method fails to meet the challenge. P1: CoIL , P2: Cao et al. , P3: Nguyen et al. , P4: ILEED , P5: Belkhale et al. , P6: Du et al. , P7: TOSIL , P8: MILES , P9: SCIZOR , P10: Hejna et al. , P11: Ours (This Work)

proaches that emphasize mutual information estimation for data quality typically require more data-hungry estimators, or approaches that rely on multi-objective losses for filtering or data cleaning risk discarding task-critical information. Our method directly addresses these challenges by exploiting the structure of periodic tasks and using temporal models to generalize from a single expert-like instance, balancing mutual information maximization with task-specific constraints. This significantly reduces data requirements and enables learning in settings where the collection of multiple high-quality demonstrations is impractical.

II. PRELIMINARIES AND PROBLEM FORMULATION

We aim to learn high-curvature periodic motion tasks, such as wiping, weaving, and pick-and-place tasks on an n-DoF robotic manipulator (in this study, we use the 7-DoF Franka Research 3 Panda robot [11]), where the end effector (EE) exhibits structured rhythmic behaviors in Cartesian space X . Mathematically, the input to our pipeline is a collection of M demonstration trajectories:

$$E = \left\{ \{x_i(t_k)\}_{k=1}^N \right\}_{i=1}^M, \quad x_i(t_k) \in \mathbb{R}^d,$$

where $x_i(t_k)$ denotes the Cartesian position of the end effector at time t_k in the i -th demonstration, and each trajectory contains N samples. The goal is to learn the continuous-time motion dynamics, $\dot{x} = f(x)$, $x \in X \subset \mathbb{R}^d$, where d can be 2 or 3 depending on the 2D or 3D Cartesian space for EE, $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ defines a nominal vector field that governs the end effector dynamics. This vector field corresponds to the Cartesian velocity of the EE, which is obtained via forward kinematics from the robot's joint velocities, tracked using a low-level impedance-based controller. In practice, however, the demonstration set E can include noisy, suboptimal, or inconsistent trajectories that degrade the learned policy \hat{f}_E obtained through standard Imitation Learning (IL) methods [12], [13], [14]. Our goal is to synthesize a corrected dataset E_{cleaned} such that $E_{\text{cleaned}} = \left\{ \{\hat{x}_i(t_k)\}_{k=1}^N \right\}_{i=1}^M$, and the resulting learned policy $\hat{f}_{E_{\text{cleaned}}}$, generates trajectories through roll out that maximize the mutual information between random variables for the state sequence $S(x_i) = \{x_i(t_k)\}_{k=1}^N$ and action sequence $A(x_i) = \{a_i(t_k)\}_{k=1}^N$, $a_i(t_k) = \frac{x_i(t_{k+1}) - x_i(t_k)}{t_{k+1} - t_k}$, $a_i(t_N) = 0$, denoted as $\mathcal{I}(S; A)$:

$$\mathcal{I}(S; A) = \mathcal{H}_S + \mathcal{H}_A - \mathcal{H}_{(S,A)}, \quad (1)$$

where entropy \mathcal{H}_X is the measure of uncertainty for the random variable X . Furthermore, we need to minimize the user-defined task intention-based loss functions $L(x_i)$, which characterizes the required task features such as symmetry, task-specific shape, smoothness, and closed-loop behavior more effectively than those generated by \hat{f}_E :

$$\mathbb{E}_{x \sim \text{Rollout}(\hat{f}_{\text{cleaned}})}[U(x)] > \mathbb{E}_{x \sim \text{Rollout}(\hat{f}_E)}[U(x)], \quad (2)$$

where $U(x) = (1 - \beta)\mathcal{I}(S(x); A(x)) - \beta L(x)$ is the *demo utility score*, $\beta \in (0, 1)$ is the demonstrator ability estimate (for expert $\beta = 1$), and $x \sim \text{Rollout}(\hat{f})$ denotes that the state trajectory $x(t)$, which is a solution of $\dot{x} = \hat{f}(x)$ starting from an initial condition. A higher demo utility score $U(x)$ improves behavior cloning by preserving informative state-action structure via mutual information (as shown by [10]) while enforcing task consistency through $L(x)$, thus reducing reliance on large datasets.

Unlike prior works, we do not assume access to a large set of high-quality expert demonstrations. Instead, we assume that the dataset contains a single expert-like trajectory $e^* \in E$. Using this, we construct a compact representation of the task via periodic Dynamic Movement Primitives (DMPs). A trajectory sequence-to-sequence correction model (e.g., LSTM) then maps noisy trajectories toward this expert-like reference, enabling policy learning from weak supervision with improved robustness and generalization. Moving forward, we will demonstrate the method for cleaning a dataset where EE is in a 2D Euclidean space, which can be easily generalized for a 3D Euclidean space for the end effector.

III. PROPOSED APPROACH

This section outlines a novel, data-efficient hierarchical approach to robot motion planning (as shown in Fig. 2) that integrates offline learning from suboptimal demonstrations.

A. Demo Utility Score: An Information-Aware Multi-objective Scoring Function

The raw set of demonstrations is often collected from teleoperation or crowdsourcing (regardless of expertise). These naive recordings are frequently of low and inconsistent quality, often presenting as noisy, inaccurate, or even trash data due to errors, incomplete tasks, or inefficient movements (as illustrated in Fig. 4). Consequently, a correction strategy is designed to transform this raw input into a reliable input suitable for live robot operation using IL, using only a few demonstrations. Each demonstration trajectory for tasks is

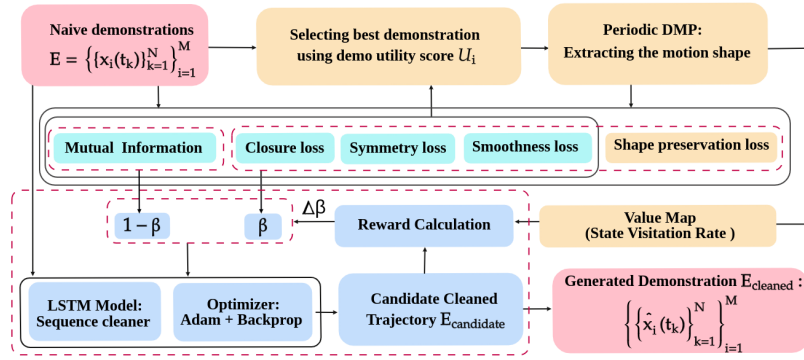


Fig. 2: Control flow diagram of the proposed approach. Pink blocks represent the set of demonstrations, yellow blocks represent the first step in the strategy to extract the required task motion pattern and value map on task space, blue blocks represent the second step in the strategy to correct the naive demonstrations with adaptive β , and cyan blocks represent the mutual information and multi-objective loss functions from task intention.

provided with general prior demonstrator preferences applied during the drawing of imperfect trajectories, such as the rotation and scaling parameters, θ in degrees and $\rho \in \mathbb{R}$, respectively. We employ a multi-objective scoring mechanism derived for task-intentional prompt, given by Do's and Don'ts of the task, easily accessible from the task manual. This quantitative evaluation allows us to prioritize the desired actions for a given state and correct the demonstrations of the entire raw dataset.

Demo utility score (U): $U(x_i)$ is calculated for each raw demonstration trajectory x_i in E . This score quantifies the quality of each demonstration based on mutual information between states and actions, as well as other important kinematic and geometric characteristics of the trajectory itself. Using the intrinsic characteristics of the sequence of pairs $(x_i(t_k), x_i(t_{k+1}))$, $\forall k \geq 1$, we compute mutual information and various loss components; a higher mutual information and a lower loss value denote a more desirable characteristic. The demo utility score $U(x_i)$, is defined as a convex combination of mutual information $\mathcal{I}(S; A)$ and the multi-objective composite loss function $L(x_i)$ for N individual loss components, given the demonstrator i 's ability estimate $\beta_i \in (0, 1)$:

$$U(x_i) = (1 - \beta_i)\mathcal{I}(S(x_i); A(x_i)) - \beta_i L(x_i), \quad (3)$$

$$L(x_i) = \sum_j w_j \cdot L_j(x_i), \forall j \in \{1, 2, \dots, N\},$$

where $L_j(x_i)$ is the j -th loss component calculated for the raw demonstration x_i , and $w_j \in [0, 1]$ is its corresponding weight.

Mutual information $\mathcal{I}(S; A)$ between the state sequence $S(x_i)$ and the action sequence $A(x_i)$ can be estimated using the Kraskov–Stögbauer–Grassberger (KSG) estimator [15] without explicit density modeling. The key steps are the following.

- For each sample, determine the distance to its k -nearest neighbor in the joint (S, A) space, which defines a local neighborhood radius.
- Count separately how many neighbors fall within this

radius in the marginal spaces S and A , which act as proxies for the marginal densities.

- Combine these counts using the digamma-based KSG formula, shown in Eq. (30) of [15], yielding an estimate of $\mathcal{I}(S; A) = \mathcal{H}_{S(x_i)} + \mathcal{H}_{A(x_i)} - \mathcal{H}_{(S(x_i), A(x_i))}$.

Intuitively, when the state sequence $S(x_i)$ and action sequence $A(x_i)$ are strongly dependent, the distance to the k -th nearest neighbor in the joint space (S, A) becomes smaller, indicating that the state–action pairs cluster more closely together. At the same time, the number of neighbors within the same joint radius is larger when viewed in the marginal spaces of S and A . In imitation learning, maximizing $\mathcal{I}(S; A)$ therefore favors demonstrations where states and actions are consistently related, leading to policies that are more robust under noisy or scarce data.

Apart from mutual information, the specific formulations for these individual loss components, based on usual task sanity checks, are given as follows, assuming $x_i = \{p_{ik}, q_{ik}\}_{k=1}^N$ represents the trajectory point as a vector of 2D coordinates (e.g. $(p_{ik}, q_{ik}) \in \mathbb{R}^2$).

- **Symmetry Loss (L_1):** It quantifies the symmetry of a trajectory with respect to the y-axis. Let C_{ix} be the x-coordinate of the centroid of the trajectory, and $C_{ix} = \frac{1}{N} \sum_{k=1}^N p_{ik}$.

$$L_1(x_i) = \frac{1}{N} \sum_{k=1}^N \min_{j \in \{1, \dots, N\}} \lambda_i, \quad (4)$$

$$\lambda_i = \|(2C_{ix} - p_{ik}, q_{ik}) - (p_{ij}, q_{ij})\|_2,$$

where $\|\cdot\|_2$ denotes the Euclidean norm. The given loss can be easily changed for symmetry across a given axis by aligning the y-axis along the required axis.

- **Closure Loss (L_2):** It quantifies the distance between the start and end point of a trajectory, which creates a closed loop.

$$L_2(x_i) = \|(p_{i0}, q_{i0}) - (p_{iN}, q_{iN})\|_2, \quad (5)$$

where $(p_{i0}, q_{i0}), (p_{iN}, q_{iN})$ is the start and end point of the given trajectory x_i .




Fig. 3: Flow diagram depicting the procedure for computing the demo utility score based on the demonstrations.

- *Smoothness Loss* (L_3): The smoothness of a trajectory is evaluated using a weighted sum of the squared norms of its first, second, and third-order finite differences, which correspond to velocity, acceleration, and jerk, respectively. Let the trajectory be represented as a sequence of N points $\{x_i(t_k)\}_{k=1}^N$ in \mathbb{R}^2 . Define D , D^2 , and D^3 as first, second, and third-order difference matrices. The smoothness loss is computed as:

$$L_3(x_i) = \alpha_1 \|D\Gamma_i\|_2^2 + \alpha_2 \|D^2\Gamma_i\|_2^2 + \alpha_3 \|D^3\Gamma_i\|_2^2, \quad (6)$$

where $\Gamma_i \in \mathbb{R}^{N \times 2}$ is the trajectory, and α_1 , α_2 , and α_3 are weights (e.g. $\alpha_1 = 0.2$, $\alpha_2 = 0.3$, $\alpha_3 = 0.5$) that balance the contributions of velocity, acceleration, and jerk to the total loss.

Using min-max normalization, we rescale mutual information $\mathcal{I}(S; A)$ as $\tilde{\mathcal{I}}$ and each loss component $-L_j(x_i)$ as $(1 - \tilde{L}_j(x_i))$ to the $[0, 1]$ interval by subtracting the minimum value and dividing by the range, with small δ , for any metric I as

$$\tilde{I}_j = \begin{cases} \frac{I_j - \min_j I_j}{\max_j I_j - \min_j I_j}, & \text{if } \max_j I_j - \min_j I_j > \delta, \\ 0, & \text{otherwise,} \end{cases} \quad (7)$$

then the normalized demo utility score for the trajectory x_i , $\tilde{U}_i(x_i) \in [0, 1]$ is computed as

$$\tilde{U}_i(x_i) = (1 - \beta_i)\tilde{\mathcal{I}}(S(x_i); A(x_i)) + \beta_i \sum_j w_j (1 - \tilde{L}_j(x_i)). \quad (8)$$

For the initial demonstration selection for learning without expert trajectories, our specific formulation of task intention score $1 - \tilde{L}(x_i)$ uses the following combination of losses:

$$1 - \tilde{L}(x_i) = \sum_{j=1}^3 w_j \cdot (1 - \tilde{L}_j(x_i)),$$

where $\tilde{L}_1, \tilde{L}_2, \tilde{L}_3$ are the normalized symmetry, closure, and smoothness loss, respectively, and the weights ($w_1, w_2, w_3 \in [0, 1], \sum_j w_j = 1$) signify the relative priority of each characteristic (shown in Fig. 3), allowing for task-specific tuning. It is crucial to note that additional task-specific loss components can be incorporated into $\tilde{L}(x_i)$, guided by user preferences.

B. Step 1: Canonical Motion Extraction with Periodic DMP

In this step, periodic DMPs [16] are used to capture a smooth and temporally consistent motion pattern from a single demonstration, rather than denoising the entire

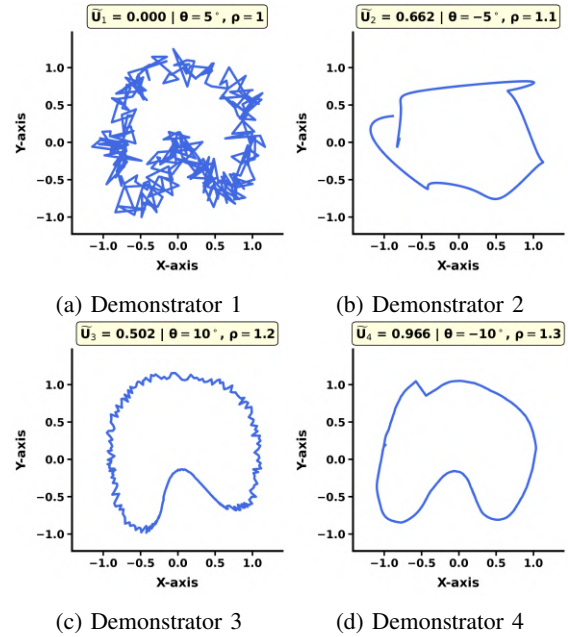


Fig. 4: Illustration of the learning setup where demonstrator or teacher trajectories (a),(b),(c),(d) are presented with associated rotation and scaling preferences (θ and ρ) to learn the task. The x and y axes represent the 2D spatial coordinates.

dataset. We have the core assumption that within the set of naive demonstrations, there exists at least one trajectory that reflects the best spirit of the task. To implement the shape extraction procedure, we follow the steps below:

- 1) The expert-like trajectory is selected using the demo utility score \tilde{U}_i (as in (8)), defined over the raw demonstration set E (shown in Fig. 4), as $x^{\text{best}}(t) := \arg \max_{x_i \in E} \tilde{U}_i(x_i)$.
- 2) After learning from the best raw demonstration ($x^{\text{best}}(t)$), the Periodic DMP generates a canonical reference for the desired motion, denoted as $x^d(t)$.
- 3) Rollouts from the learned Periodic DMP are then used to estimate a state visitation value map $V_E(s) : \mathbb{R}^2 \rightarrow [0, 1]$ (as shown in Fig. 5) over the task space, by estimating state visitation density using the nonparametric kernel density estimation (KDE) technique (described in [17]). Formally, for a set of rollout states $\{x_i(t_k)\}_{k=1}^N$ in a 2D task space, the KDE estimate at a point $x_i(t_k)$ is given by $\hat{p}(x_i(t_k)) = \frac{1}{Nh^2} \sum_{i=1}^N K\left(\frac{x - x_i(t_k)}{h}\right)$, where K is a kernel function (e.g., Gaussian) and h is the bandwidth controlling smoothness. The resulting $\hat{p}(x_i)$ serves as the state visitation value map. This approach is advantageous because it does not require discretizing the task space arbitrarily. This value map guides the reward $R_\pi \in \mathbb{R}$ calculation for the learned policy $\pi(s) : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, defining the map from state to action, in the refinement stage, required for adapting the demonstrator's ability factor β during training.
- 4) To preserve the variability present in the dataset, the shape trajectory $x^d(t)$ is inverse transformed as per the

rotation and scaling parameter, derived from the best trajectory $x^{\text{best}}(t)$'s demonstrator preference, in order to obtain the global shape reference $x^g(t)$ with $\theta = 0, \rho = 1$, i.e., no demonstrator bias.

- 5) For every raw demo $x_i(t)$, the corresponding shape reference is generated from $x^g(t)$, using its respective demonstrator's rotation and scaling parameters θ, ρ , giving $x_i^{\text{aligned}}(t)$.
- 6) Calculate the loss $L_4(x_i)$ as the distance between the respective demonstrator canonical trajectory and raw demo:

$$L_4(x_i) = \frac{1}{M} \sum_{i=1}^M \sum_{k=1}^N \left\| x_i^{\text{aligned}}(t_k) - x_i(t_k) \right\|_2. \quad (9)$$

A key result of this stage is the shape preservation score $1 - \tilde{L}_4(x_i)$ (normalized as in (7)), quantifying the accuracy with which the canonical representation reflects the fundamental shape shared among the set of imperfect demonstrations. This loss, along with other intentional loss components such as (4),(5),(6), will then be used in the subsequent learning based trajectory refinement stage for optimization as

$$1 - \tilde{L}(x_i) = \sum_{j=1}^4 w_j \cdot (1 - \tilde{L}_j(x_i)). \quad (10)$$

C. Step 2: Refinement with Demo utility score

We introduce a Long Short-Term Memory (LSTM) model for refinement and correction. LSTMs are particularly well suited for capturing long-term temporal dependencies [18], enabling the model to learn context-aware corrections essential for complex movements. The LSTM is trained as a sequence-to-sequence model that maps a noisy demonstration trajectory $x_i(t) \in E$ to a denoised and structurally refined output trajectory $\hat{x}_i(t)$. The training objective for the network is to learn this mapping by maximizing the demo utility score \tilde{U}_i defined in (8), using (4), (5), (6), and (9). This formulation enables the LSTM to learn corrections that align the raw input trajectories with the task-intended behavior without requiring explicit supervision through human interventions.

During training, the demonstrator ability factor β is adapted based on the rollout trajectory generated by the LSTM policy $\pi(s)$. Specifically, the reward R_π for a rollout is computed by summing the state visitation probabilities

$V_E(s)$ of all states visited by the trajectory:

$$R_\pi = \sum_{s \in \text{Rollout}(\pi)} V_E(s). \quad (11)$$

The change in the mean reward $\hat{R}_\pi^{(k)}$ per epoch k is then used to update β_i via a tanh function, for increased sensitivity via $\gamma \in \mathbb{R}$ even with small change in reward:

$$\begin{aligned} \Delta \hat{R}_\pi^{(k)} &= \gamma (\hat{R}_\pi^{(k)} - \hat{R}_\pi^{(k-1)}), \\ \beta_i^{(k+1)} &= \text{clip}(\beta_i^{(k)} + \tanh(\Delta \hat{R}_\pi^{(k)}), 0, 1), \end{aligned}$$

If the reward decreases, β_i is reduced, increasing the emphasis on maximizing mutual information $\mathcal{I}(S; A)$ to guide learning towards the embedded optimal policy. Conversely, if the reward increases, β_i is raised, placing more weight on minimizing the task intention loss to ensure adherence to task objectives. In general, this adaptive principle for tuning β_i allows the LSTM to progressively enhance the behavior of the embedded demonstrator, required to clean the raw demonstrations. Even when the inputs are significantly distorted, the model is able to produce outputs that closely resemble the expert demonstration in shape, periodicity, and symmetry (as conceptually shown in Fig. 2 and illustrated by results in Fig. 6). This two-stage refinement process, illustrated in Algorithm 1, leverages the generalization capabilities of DMPs for canonical motion extraction and the expressiveness of LSTMs for fine-grained refinement, which provides a robust framework for learning from a small number of non-expert or degraded demonstrations. The final refined trajectory, denoted as $E_{\text{cleaned}} = \{\hat{x}(t_k)\}_{k=1}^N$, where $\hat{x}(t_k)$ represents the refined state at discrete time t_k (conceptually derived from an integral such as $\hat{x}(t) = \int f(\cdot) dt$), then can be used as a high-quality target for any subsequent control or tracking systems such as the Neural ODE based low-level controller design as in [19], which is validated with experiments in Section V.

IV. COMPARISON AND DISCUSSIONS

We have four raw demonstrations, denoted as dataset D_1 , for wiping task (WP), pick-and-place task (PnP) and weaving task (WV) as in Fig. 7(a), 7(b), 7(c) and corresponding cleaned demonstrations E_{cleaned} as in Fig. 7(d), 7(e), 7(f), respectively, where $N = 300, M = 4$, with rotation parameters $\{\theta_i\}_{i=1}^4 = \{5^\circ, -5^\circ, 10^\circ, -10^\circ\}$ and $\{\rho_i\}_{i=1}^4 = \{1, 1.1, 1.2, 1.3\}$. These variations in θ and ρ introduce demonstrator preferences, allowing us to evaluate the robustness of our framework in recovering the underlying trajectory shapes from distorted data. We also have another 20 raw demonstrations, denoted as dataset D_2 , especially to validate the IL methods, which are data-hungry. These were collected from the same four instructors, where each instructor provided five demonstrations under the same preference parameters θ_i and ρ_i .

We then train multiple imitation learning methods, including Behavioral Cloning (BC) [20], Trajectory-based BC (Traj-BC) [21], Imitation Learning by Estimating Expertise of Demonstrators (ILEED) [4], and Neural ODE (NODE)




Fig. 5: State visitation value maps from rollouts of the learned periodic DMP model across tasks.




Fig. 6: Teacher demonstrations $\{x_i(t_k)\}_{i=1}^4$ shown as solid blue lines in (a), their corresponding shape reference trajectories $\{x_i^r(t_k)\}_{i=1}^4$, denoted as target trajectory, shown as dotted red line, and the resulting cleaned outputs $\{\hat{x}_i(t_k)\}_{i=1}^4$ shown as solid blue lines in (b). Each demonstration is refined using its shape reference to produce a consistent and structured set of cleaned demonstrations.

Algorithm 1 Two-stage Pipeline for Data Cleaning

- 1: **Input:** Raw demonstration set $E = \{\{x_i(t_k)\}_{k=1}^N\}_{i=1}^M$
- 2: **Initialize:** Demonstrator preference parameters ρ, θ , demonstrator ability factor β_i for all demonstrators, LSTM network weights W , learning rate η .
- 3: Compute demo utility scores $\tilde{U}(x_i)$ for all $x_i \in E$

Step 1: Canonical Motion Extraction with Periodic DMP

- 4: Select best demo $x^{\text{best}}(t) \leftarrow \arg \max_{x_i \in E} \tilde{U}(x_i)$
- 5: Extract shape reference $x^d(t)$ using periodic DMP trained on $x^{\text{best}}(t)$
- 6: Generate canonical shape $x^g(t)$ by removing $x^d(t)$ demonstrator's rotation θ and scale ρ bias, derived from $x^{\text{best}}(t)$ demonstrator.
- 7: Generate shape reference trajectory $x_i^{\text{aligned}}(t)$ for each demonstrator by rotating and scaling $x^g(t)$ by parameter ρ and θ , respectively.
- 8: Compute shape preservation loss $L_4(x_i)$
- 9: Compute $(1 - \tilde{L}(x_i)) \leftarrow \sum_{j=1}^4 w_j \cdot (1 - \tilde{L}_j(x_i))$

Step 2: Refinement using LSTM

- 10: **for** each demonstration $x_i \in E$ **do**
 - 11: **for** each training epoch $k = 1$ to K **do**
 - 12: **for** each LSTM rollout using policy π **do**
 - 13: Compute task intention score $1 - \tilde{L}(x_i)$
 - 14: Compute demo utility score \tilde{U}_i
 - 15: $W^{(k+1)} \leftarrow W^{(k)} - \eta \nabla_{\theta}(-\tilde{U}(x_i))$
 - 16: Compute reward $R_{\pi} \leftarrow \sum_{s \in \text{Rollout}(\pi)} V_E(s)$
 - 17: **end for**
 - 18: Update $\beta_i^{(k+1)} \leftarrow \text{clip}(\beta_i^{(k)} + \tanh(\gamma \Delta \hat{R}_{\pi}), 0, 1)$
 - 19: **end for**
 - 20: **Output:** LSTM rollout $\hat{x}_i(t)$
 - 21: **end for**
 - 22: **Output:** $E_{\text{cleaned}} = \{\{\hat{x}_i(t_k)\}_{k=1}^N\}_{i=1}^M$
-

[19] on both cleaned and noisy datasets for D_1 as well as D_2 . The resulting outputs for the dataset D_1 , corresponding to 50 rollouts of each trained models, are as shown in Fig. 8, Fig. 9, and Fig. 10 which show that the standard IL methods

perform well especially with respect to the loss L_4 for shape preservation with dataset E_{cleaned} , demonstrating the effectiveness of our approach for complex imitation learning tasks like periodic motions.

It also validates the statistical advantage by evaluating the performance of the rollout trajectories against the reward R_{π} (as computed in (11)), which is higher for rollout states with higher value in the map V_E , estimated by the single expert-like trajectory, as shown by the bar chart in Fig. 11. Our method offers the highest statistical advantage in behavior cloning when integrated with ILEED for dataset D_1 , with increase in mean normalized reward for wiping task by 0.88, pick-and-place task by 0.418, and weaving task by 0.084, as it combines knowledge of demonstrator skill with a corrective refinement mechanism that effectively "skills up" the demonstrator, reducing reliance on knowledge of the complete state space with large dataset. This synergy allows data-hungry ILEED to achieve significantly improved performance even with a small dataset, containing as few as four trajectories, highlighting the robustness and data-efficiency of our approach. Similarly, a statistical advantage is observed for dataset D_2 as well.

The Neural ODE (NODE) model demonstrates consistently stronger performance across all tasks when trained on the cleaned dataset compared to the noisy dataset. This improvement can be attributed to the higher mean demo utility score of the cleaned dataset, preserving high mutual information between the state and action sequences in the cleaned data, which allows the model to better encode the underlying dynamics. This effect is particularly important in the weaving task, where capturing long-range dependencies and the structure of the two conjoined loops is crucial for accurate trajectory generation.

BC and Traj-BC models achieve good performance on the wiping task, benefiting from task intention scores that promote symmetry and smooth closure despite limited data. However, their strong data dependence prevents them from learning the more complex pick-and-place and weaving tasks with only four demonstrations, as evidenced by the high variance in normalized rewards. It is also unable to learn the task behavior with 20 demos, but the variance in rewards is reduced. However, the rollout trajectories for BC and Traj-BC models trained on the cleaned dataset are able to preserve

the shape for the required task implementation as compared to the models trained on the noisy dataset, especially in the case of the weaving task.

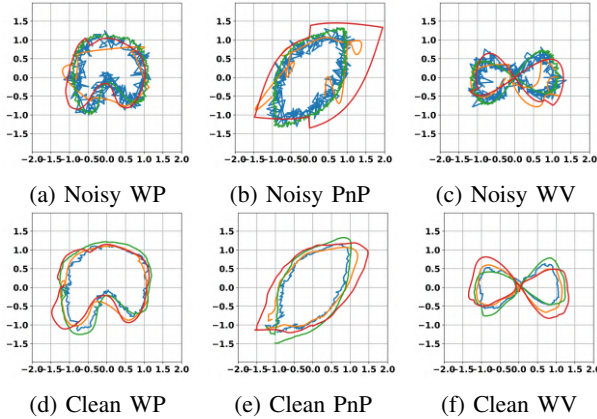


Fig. 7: Comparison of noisy versus cleaned trajectories (D_1) for Wiping task (WP), Pick-and-place task (PnP) and Weaving task (WV).

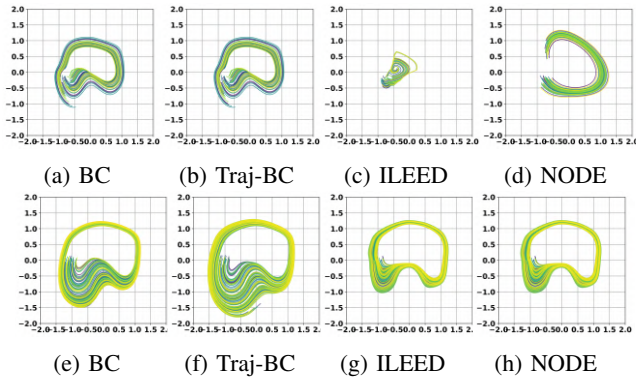


Fig. 8: Comparison of rollouts across four IL models trained on noisy D_1 (top row) and cleaned (bottom row) demonstrations for wiping task.

V. EXPERIMENTAL RESULTS

To validate our method, we performed experiments on the Franka Research 3 Panda robot using a hardware setup as shown in Fig. 12. For hardware video visit <https://youtu.be/sEGpXu0vF7Q>. The 2D demonstration trajectories in $E, E_{cleaned}$ are padded with a height of 0.3 unit along the z-axis and centered about the point (0.35, 0) for the 3D trajectories required for the wiping task. Similarly, the demonstration trajectories were aligned in 3D space for the pick-and-place task. Since the NODE framework did not learn the task shape for both $E, E_{cleaned}$ (cf. Fig. 10) with the weaving task, we excluded it in the final execution. In the first scenario, we learned the nominal trajectory using *only 4 naive demonstrations for the wiping task* as shown in Fig. 4, with the help of the NODE framework [19]. The resulting motion was inconsistent and failed to preserve the desired specifications as shown in

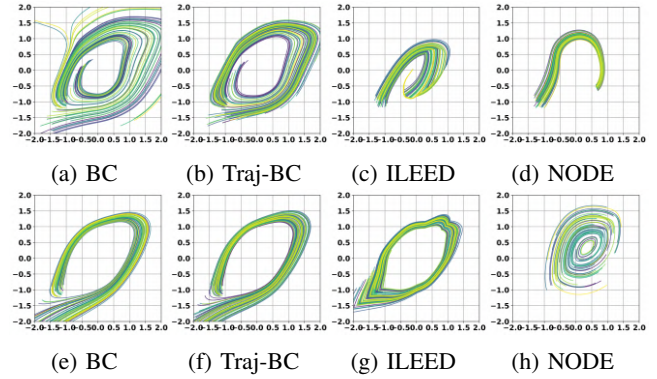


Fig. 9: Comparison of rollouts across four IL models trained on noisy D_1 (top row) and cleaned (bottom row) demonstrations for pick-and-place task.

Fig. 12a, 12c. In contrast, when using our corrected set of demonstrations $E_{cleaned}$, the learned trajectory aligned closely with the intended expert-like drawing. The corrected trajectory retained the essential curvature, periodicity, and closure properties, demonstrating the model’s ability to robustly recover trajectories as shown in Fig. 12b, 12d.

VI. CONCLUSION AND LIMITATION

This work introduces a novel, data-efficient pipeline that overcomes the long-standing reliance on extensive, high-quality datasets for imitation learning. By leveraging a two-stage process, first extracting a canonical motion shape from a single expert-like demonstration via Periodic Dynamic Movement Primitives (DMPs), then refining the entire set of noisy demonstrations with a multi-objective LSTM, our framework successfully recovers a clean, expert-quality dataset. We have demonstrated that this data-centric approach enables standard imitation learning models to achieve robust, expert-like robotic performance on complex periodic tasks, even when trained on as few as four demonstrations. Even with 20 cleaned demos for dataset D_2 , our approach provides better shape-preserving actions as compared to noisy demos.

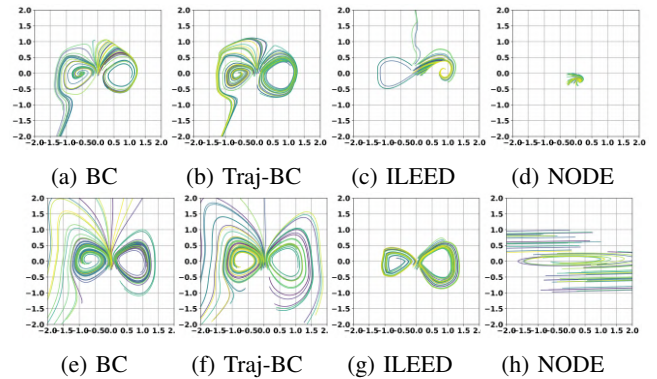


Fig. 10: Comparison of rollouts across four IL models trained on noisy D_1 (top row) and cleaned (bottom row) demonstrations for the weaving task.

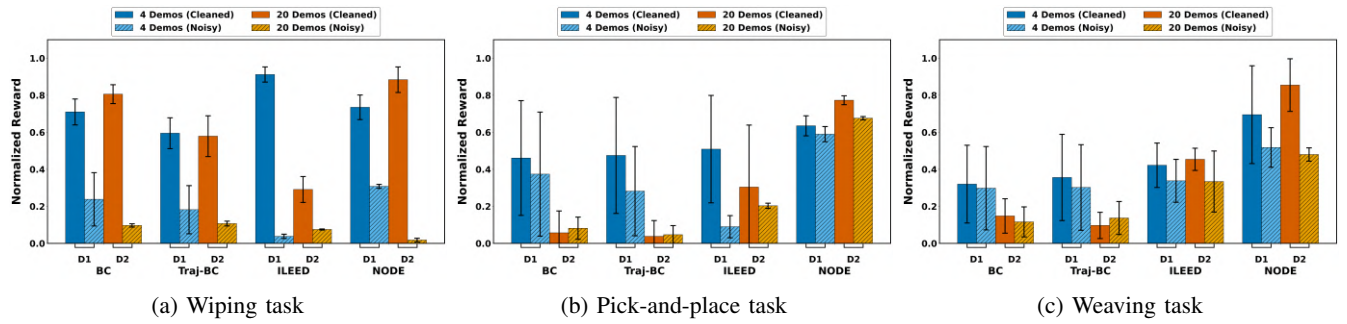


Fig. 11: Comparison of mean reward and reward variability (standard deviation) for rollout trajectories generated by models trained on noisy versus cleaned demonstration datasets D_1, D_2 .

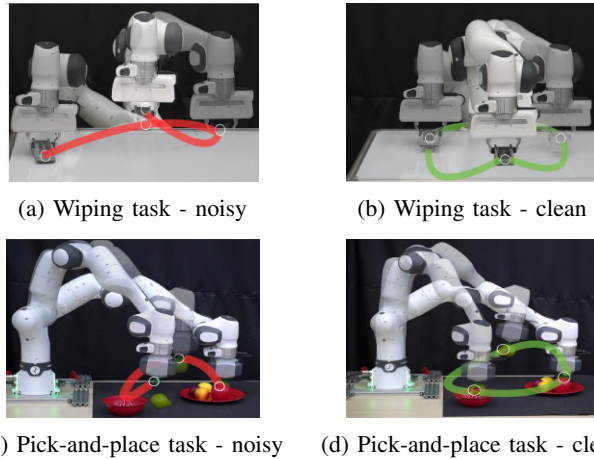


Fig. 12: Experimental results with expert-like behavior using imperfect demonstrations. Here “noisy” and “clean” mean rollouts are from IL models trained on dataset E and E_{cleaned} , respectively. For video visit <https://youtu.be/sEGpXu0vF7Q>.

A key limitation, however, is the assumption that at least one expert-like demonstration exists within the initial dataset to serve as a reference. Despite this, our results across diverse domains: wiping, weaving, and pick-and-place, highlight that data quality can be manufactured and that a single reference of expertise is a powerful substitute for large-scale data collection. This work presents a significant step toward making imitation learning more practical and accessible for real-world robotic applications where obtaining perfect, abundant data is a major bottleneck.

REFERENCES

- [1] G. Papagiannis and E. Johns, “MILES: making imitation learning easy with self-supervision,” in *8th Conference on Robot Learning*, pp. 810–829, 2025.
- [2] H. Yu, Y. Zhu, T. Li, and W. Wu, “Co-imitation learning without expert demonstration,” *arXiv preprint arXiv:2103.14823*, 2021.
- [3] J. Su, Y. Li, Z. Wu, Z. Chen, S. Li, and R. Li, “Task-oriented self-imitation learning for robotic autonomous skill acquisition,” *International Journal of Humanoid Robotics*, vol. 21, no. 02, p. 2450001, 2024.
- [4] M. Beliaev, A. Shih, S. Ermon, D. Sadigh, and R. Pedarsani, “Imitation learning by estimating expertise of demonstrators,” in *International Conference on Machine Learning*, pp. 1732–1748, PMLR, 2022.
- [5] Z. Cao and D. Sadigh, “Learning from imperfect demonstrations from agents with varying dynamics,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5231–5238, 2021.
- [6] H. H. Nguyen, A. Baisero, D. Wang, C. Amato, and R. Platt, “Leveraging fully observable policies for learning under partial observability,” in *Conference on Robot Learning*, pp. 1673–1683, PMLR, 2023.
- [7] S. Belkhale, Y. Cui, and D. Sadigh, “Data quality in imitation learning,” *Advances in neural information processing systems*, vol. 36, pp. 80375–80395, 2023.
- [8] M. Du, S. Nair, D. Sadigh, and C. Finn, “Behavior retrieval: Few-shot imitation learning by querying unlabeled datasets,” *arXiv preprint arXiv:2304.08742*, 2023.
- [9] Y. Zhang, Y. Xie, H. Liu, R. Shah, M. Wan, L. Fan, and Y. Zhu, “Scizor: A self-supervised approach to data curation for large-scale imitation learning,” in *Workshop on Making Sense of Data in Robotics: Composition, Curation, and Interpretability at Scale at CoRL 2025*.
- [10] J. Hejna, S. Mirchandani, A. Balakrishna, A. Xie, A. Wahid, J. Tompson, P. Sanketi, D. Shah, C. Devin, and D. Sadigh, “Robot data curation with mutual information estimators,” *arXiv preprint arXiv:2502.08623*, 2025.
- [11] S. Haddadin, “The Franka Emika robot: A standard platform in robotics research,” *IEEE Robotics & Automation Magazine*, 2024.
- [12] K. Eder, C. Harper, and U. Leonards, “Towards the safety of human-in-the-loop robotics: Challenges and opportunities for safety assurance of robotic co-workers,” in *23rd IEEE International Symposium on Robot and Human Interactive Communication*, pp. 660–665, 2014.
- [13] Z. Yuan, A. W. Hall, S. Zhou, L. Brunke, M. Greeff, J. Panerati, and A. P. Schoellig, “Safe-control-gym: A unified benchmark suite for safe learning-based control and reinforcement learning in robotics,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11142–11149, 2022.
- [14] D. Totsila, K. Chatzilygeroudis, V. Modugno, D. Hadjiveličkov, and D. Kanoulas, “Sensorimotor learning with stability guarantees via autonomous neural dynamic policies,” *IEEE Robotics and Automation Letters*, 2025.
- [15] A. Kraskov, H. Stögbauer, and P. Grassberger, “Estimating mutual information,” *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, vol. 69, no. 6, p. 066138, 2004.
- [16] A. Ude, A. Gams, T. Asfour, and J. Morimoto, “Task-specific generalization of discrete and periodic dynamic movement primitives,” *IEEE Transactions on Robotics*, vol. 26, no. 5, pp. 800–815, 2010.
- [17] B. W. Silverman, *Density estimation for statistics and data analysis*. Routledge, 2018.
- [18] Y. Duan, Y. L. V., and F.-Y. Wang, “Travel time prediction with LSTM neural network,” in *IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1053–1058, 2016.
- [19] F. Nawaz, T. Li, N. Matni, and N. Figueroa, “Learning complex motion plans using neural ODEs with safety and stability guarantees,” in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 17216–17222, 2024.
- [20] F. Torabi, G. Warnell, and P. Stone, “Behavioral cloning from observation,” in *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pp. 4950–4957, 2018.
- [21] C. A. Hepburn and G. Montana, “Model-based trajectory stitching for improved behavioural cloning and its applications,” *Machine Learning*, vol. 113, no. 2, pp. 647–674, 2024.