

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 HIERARCHICAL VALUE-DECOMPOSED OFFLINE REIN- FORCEMENT LEARNING FOR WHOLE-BODY CONTROL

Anonymous authors

Paper under double-blind review

ABSTRACT

Scaling imitation learning to high degrees-of-freedom (DoF) whole-body robots is fundamentally limited by the increased DoF and the non-stationary observation transition. We argue that the core bottleneck is paradigmatic: real-world supervision for whole-body control is inherently imperfect, yet most methods assume expert data. To overcome this, we propose **HVD** (Hierarchical Value-Decomposed Offline Reinforcement Learning), an offline reinforcement learning framework that learns effective policies directly from suboptimal, reward-labeled trajectories. HVD structures the value function along the robot’s kinematic hierarchy and over temporal chunks, enabling precise credit assignment in long-horizon, high-dimensional tasks. Built on a Transformer-based architecture, HVD supports *multi-modal* and *multi-task* learning, allowing flexible integration of diverse sensory inputs. To enable realistic evaluation and training, we further introduce **WB-50**, a 50-hour dataset of teleoperated and policy rollout trajectories annotated with rewards and preserving natural imperfections — including partial successes, corrections, and failures. Experiments show HVD significantly outperforms existing baselines in success rate across complex whole-body tasks. Our results suggest effective policy learning for high-DoF systems can emerge not from perfect demonstrations, but from structured learning over realistic, imperfect data.

1 INTRODUCTION

Imitation learning has become a central paradigm for developing robotic policies, enabling robots to acquire complex skills by learning from expert demonstrations (Pomerleau, 1991; Ross et al., 2011; Brantley et al., 2020). Recent advances, such as diffusion-based policies for robust action generation (Chi et al., 2023; Ze et al., 2024; Liu et al., 2024) and Vision-Language-Action (VLA) models for instruction following (Kim et al., 2024; Black et al., 2024; Liu et al., 2024; Hu et al., 2024; Intelligence et al., 2025), extend imitation learning into unified perception-reasoning-control frameworks, establishing it as a foundation for versatile, general-purpose robotic systems.

However, scaling imitation learning from manipulators to whole-body systems with high-DoF introduces fundamental challenges due to the increased DoF (Bellman, 1966; Kober et al., 2013). As the number of joints increases, both the state and action spaces expand exponentially. Moreover, unlike manipulators operating in static environments with fixed third-person cameras, whole-body robots rely on egocentric perception. Small torso or head movements can cause significant visual shifts, demanding policies robust to dynamic, self-induced viewpoint changes (Tian et al., 2025). Furthermore, teleoperating such high-dimensional systems is cognitively and physically costly, often leading to noisy, suboptimal trajectories (Zhou et al., 2023; Jiang et al., 2025). We argue that the bottleneck lies not only in the policy design but in the paradigm itself: reliance on optimal demonstrations is misaligned with the imperfect supervision available in real-world whole-body control. Thus, a key question emerges:

How can we learn desired whole-body behaviors when expert demonstrations are limited?

We observe that large volumes of non-expert data are organically generated during human teleoperation and policy rollout. These trajectories often contain partial successes, corrective maneuvers, and outright failures. Although they encode rich, scalable supervisory signals, such data have been largely overlooked in prior work. A principled way to capitalize on such data is *offline reinforcement learn-*

054 *ing* (offline RL), which learns policies entirely from fixed datasets without additional environment
 055 interaction (Levine et al., 2020). Yet, existing offline RL approaches remain ill-suited for whole-body
 056 control, as they struggle with high-dimensional action spaces, sparse reward signals, and the need to
 057 generalize across diverse tasks and multi-modal perceptual inputs (Kalashnikov et al., 2018; Kumar
 058 et al., 2022; Bhateja et al., 2023).

059 To address these issues, we propose **HVD** (Hierarchical Value-Decomposed Offline Reinforcement
 060 Learning), an offline reinforcement learning framework designed for whole-body control. HVD
 061 features a structured Q-function that jointly incorporates kinematically informed value decomposition
 062 and temporal Q-chunking (Li et al., 2025), enabling precise credit assignment and stabilizing learning
 063 across long-horizon, multi-step whole-body behaviors. Built upon a Transformer-based backbone,
 064 HVD supports multi-modal sensory inputs and scales gracefully to multi-task settings. To ground this
 065 framework in practice, we further contribute **WB-50**, a 50-hour dataset of teleoperated and policy
 066 rollout trajectories collected under realistic conditions. WB-50 is annotated with reward labels tied
 067 to subtask outcomes. It deliberately preserves natural imperfections, including partial successes,
 068 corrections, and failures, thereby offering a rich and realistic benchmark for whole-body offline RL.

069 We evaluate HVD on a diverse set of complex whole-body manipulation tasks requiring precise,
 070 coordinated control across multiple limbs. Experimental results show that HVD effectively learns
 071 from the imperfections, significantly outperforming existing baselines in success rate and task
 072 performance. Our findings show that high-DoF policy learning can be both practical and scalable—not
 073 only by pursuing sufficient perfect demonstrations, but by leveraging suboptimal data in a structured,
 074 hierarchical offline RL framework, paving the way for more accessible deployment of whole-body
 075 control in real-world robotic applications.

076 The primary contributions of this work are as follows:

- 078 • We propose HVD, an offline RL method for whole-body control via hierarchical Q-value
 079 decomposition with temporal chunking, enabling precise credit assignment in high-DoF,
 080 long-horizon tasks.
- 081 • We implement HVD using a Transformer-based architecture that supports multi-modal
 082 inputs and multi-task learning.
- 083 • We introduce WB-50, a 50-hour whole-body robotics dataset of imperfect, reward-labeled
 084 trajectories.
- 085 • We demonstrate that HVD outperforms baselines across diverse whole-body tasks and policy
 086 architectures, and validate the effectiveness of HVD on multi-task settings.

089 2 PRELIMINARIES

092 **Markov Decision Process.** We model the robot control task as a Markov Decision Process (MDP)
 093 defined by the tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{r}, H)$, where H is the horizon. The state space $\mathcal{S} = \mathcal{S}_{\text{obs}} \times \mathcal{S}_{\text{prop}}$
 094 includes egocentric observations (e.g., images, point clouds) and proprioceptive state. The action
 095 space $\mathcal{A} \subset \mathbb{R}^d$ consists of joint-level commands for a d -DoF robot. and $\mathcal{r}(s^h, a^h) \in [0, 1]$ is the
 096 reward function and $\mathcal{P}(s^{h+1}|s^h, a^h)$ characterizes the non-stationary transition function of this MDP,
 097 which is a critical assumption because the whole-body movement and the constraints of the camera's
 098 Field of View (FoV) result in partial observability, making the observation-based transitions appear
 099 highly stochastic and time-varying throughout the horizon. The goal is to learn a policy $\pi(a|s)$ that
 100 maximizes the expected return $V(\pi) = \mathbb{E}_{\pi}[\sum_{h=0}^H r(s_h, a_h)]$.

101 **Offline Reinforcement Learning** considers the problem of learning a policy from a fixed dataset
 102 $\mathcal{D} = \{(s, a, r, s')\}$ without further environment interaction. A central challenge is *distributional shift*,
 103 where the learned policy may query actions outside the support of the data, leading to erroneous
 104 value estimates and thus poor performance (Kumar et al., 2019; Levine et al., 2020; Koh et al., 2021).

105 **Implicit Diffusion Q-Learning (IDQL)** (Hansen-Estruch et al., 2023) builds on IQN Kostrikov et al.
 106 (2021), which can be viewed as an *actor-critic method* (Konda & Tsitsiklis, 1999), where the critic
 107 objective induces an implicit, behavior-regularized actor to prevent the value overestimation problem
 in offline RL. In this framework, the value function $V_{\psi}(s)$ is obtained by minimizing a convex loss

108 over dataset actions:

109

$$110 \quad V_{\psi}^*(s) = \min_{\psi} \mathbb{E}_{a \sim \mu(a|s)} [f(Q_{\theta}(s^h, a^h) - V_{\psi}(s^h))],$$

111

112 where f is chosen as an asymmetric convex function (e.g., expectiles (Kostrikov et al., 2021),
113 quantiles (Koenker & Hallock, 2001), or exponential (Beirlant et al., 1999)), determining how the
114 implicit policy π_{ϕ} deviates from the behavior policy μ . The Q-function is trained with Bellman
115 backups:

116

$$\mathcal{L}_Q(\theta) = \mathbb{E}_{(s^h, a^h, s^{h+1}) \sim \mathcal{D}} [(r(s^h, a^h) + V_{\psi}(s^{h+1}) - Q_{\theta}(s^h, a^h))^2],$$

117

118 To recover the policy, IDQL employs $\pi_{\phi}(a|s)$ trained via advantage weighted regression (Sasaki &
119 Yamashina, 2020):

120

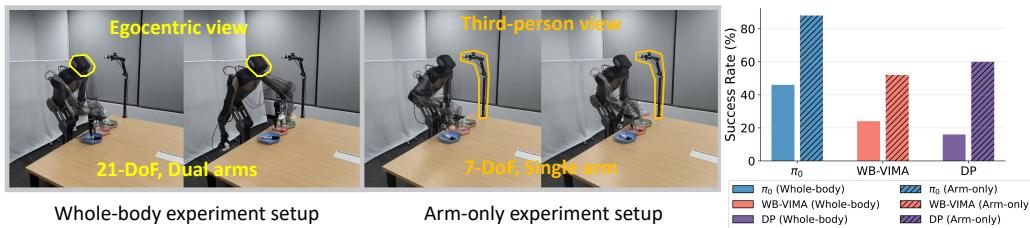
$$\mathcal{L}_{\pi}^{\text{weightbp}}(\phi) = \frac{1}{H} \sum_{h=1}^H \mathbb{E} \left[\frac{|f'(Q(s^h, a^h) - V^*(s^h)|}{|Q(s^h, a^h) - V^*(s^h)|} \left\| \epsilon - \pi_{\phi}(\sqrt{\hat{\alpha}} a^h + \sqrt{1-\hat{\alpha}} \epsilon, s^h, t) \right\| \right],$$

121

122 where $\epsilon \sim \mathcal{N}_{\text{pol}}(0, I)$ denotes Gaussian noise, t is the noising timestep, $\hat{\alpha}_t$ is the noise schedule
123 parameter in diffusion training, and $f' = \frac{\partial f}{\partial V(s)}$ denotes the derivative of f with respect to $V(s)$.
124

126 3 CHALLENGES IN WHOLE-BODY CONTROL

127 In this section, we argue that whole-body control presents two key challenges: the increased DoF
128 and non-stationary observation dynamics. We examine the impact of this phenomenon on policy
129 performance from both theoretical and empirical perspectives.



140 Figure 1: Performance comparison between 21-DoF whole-body policies and 7-DoF arm-only
141 policies trained on 50 expert demonstrations.

144 3.1 THEORETICAL ANALYSIS

145 In the context of **Behavior Cloning** (BC), theoretical studies have established that the *expert sample*
146 *complexity*, defined as the number of expert trajectories required to learn a policy with a desired level
147 of performance, scales poorly with the size of the policy set Π . This has been a focus of recent work
148 (Rajaraman et al., 2020; Tu et al., 2022; Foster et al., 2024). We begin by formally introducing the
149 policy covering number of a policy class.

150 **Definition 3.1** (Policy Covering Number). For a policy class $\Pi \subset \{\pi_h : \mathcal{X} \rightarrow \Delta(\mathcal{A})\}$ ¹, we set that
151 $\Pi' \subset \{\pi_h : \mathcal{X} \rightarrow \Delta(\mathcal{A})\}$ is an ε -cover if for all $\pi \in \Pi$, there exists $\pi' \in \Pi'$ such that for all $x \in \mathcal{X}$,
152 $a \in \mathcal{A}$, and $h \in [H]$,

$$154 \quad \log \left(\frac{\pi_h(a|x)}{\pi'_h(a|x)} \right) \leq \varepsilon. \quad (1)$$

155

156 We denote the size of the smallest such cover by $\mathcal{N}_{\text{pol}}(\Pi, \varepsilon)$.

157 We then formalize the dependence of the expert sample complexity of BC on this measure:

158

159 ¹While diffusion policies are typically implemented as implicit generative models, they theoretically induce
160 an explicit probability density function $\pi(a|x)$ via the *Probability Flow ODE* formulation (Song et al., 2021).
161 This bijective mapping ensures that the log-density $\log \pi(a|x)$ is well-defined and computable, rendering
162 diffusion policies compatible with this definition.

162 **Proposition 3.2** (Expert Sample Complexity of Behavior Cloning (Foster et al., 2024)). *For any expert*
 163 *policy $\pi^* \in \Pi$, to ensure that the suboptimality gap of the learned policy $\hat{\pi}$ satisfies $V(\pi^*) - V(\hat{\pi}) \leq \varepsilon$*
 164 *with probability at least $1 - \delta$, the number of expert trajectories n required for behavioral cloning*
 165 *needs to satisfy:*

$$166 \quad n = \tilde{\mathcal{O}} \left(\frac{H^3 \log \mathcal{N}_{pol}(\Pi, \varepsilon_\pi)}{\varepsilon^2} \right). \quad (2)$$

168 *Here, H is the task horizon, ε is the target suboptimality gap, and $\mathcal{N}_{pol}(\Pi, \varepsilon_\pi)$ denotes the ε_π -policy*
 169 *covering number of the policy class Π .*

171 This result reveals a critical quantitative limitation of BC: the expert sample complexity of BC is
 172 fundamentally tied to the log-covering number of the policy class Π . As the **complexity or size of**
 173 **the action space expands**, $\mathcal{N}_{pol}(\Pi, \varepsilon_\pi)$ **grows**, which in turn necessitates a larger number of expert
 174 trajectories n to learn an ε_π -optimal policy.

176 3.2 EMPIRICAL VALIDATION

178 Figure 1 provides an empirical comparison of arm-only and whole-body policies applied to the same
 179 task. The leftmost and middle panels show the experimental setups for both cases: the whole-body
 180 setup uses a 21-DoF robot with dual arms and an egocentric view, while the arm-only setup uses a
 181 7-DoF robot with a single arm and a third-person view. The rightmost panel compares the success
 182 rates of various policies (π_0 , WB-VIMA, and DP) with the same number of expert demonstrations
 183 (50) for both the whole-body and arm-only setups. The whole-body policies consistently achieve
 184 significantly lower success rates than their arm-only counterparts across all models. This performance
 185 gap highlights the challenges posed by high-dimensional action spaces and non-stationary transitions.

186 In summary, the increased DoF in whole-body control leads to an explosion in expert data require-
 187 ments, driven both by the complexity of DoFs and the inherent non-stationary transition.

188 4 METHOD

191 4.1 WB-50: IMPERFECT DATA FOR WHOLE-BODY RL

192 The preceding analysis shows that directly relying on expert demonstrations is impractical due to the
 193 increased DoF and instability of egocentric observations. Fortunately, in realistic settings, *non-expert*
 194 *demonstrations* are more abundant, arising naturally from teleoperation and policy rollouts (Zhou
 195 et al., 2023). To leverage this, we introduce **WB-50** as illustrated in Figure 2: a reward-labeled
 196 whole-body dataset spanning **over 50 hours**. WB-50 is intentionally composed of three distinct
 197 data sources to reflect realistic data distributions: i) **expert demonstrations** (43.7%), ii) **imperfect**
 198 **teleoperation** (14.6%), and iii) **policy rollouts** (41.7%) — the latter two comprising the majority,
 199 mirroring the scarcity of perfect supervision in practice. Every trajectory is annotated at the subtask
 200 level and labeled with discrete reward signals. More details are listed in Appendix B.

201 A direct way to leverage such data is *offline reinforcement learning* (offline RL), which enables
 202 policy learning from static and imperfect datasets (Lange et al., 2012; Levine et al., 2020). However,
 203 existing offline RL methods face fundamental limitations when applied to whole-body control. Most
 204 prior successes have been restricted to relatively low-dimensional tasks (Mandlekar et al., 2022;
 205 Sinha et al., 2022; Zhou et al., 2023), and current algorithms struggle to scale to the high degrees
 206 of freedom inherent in whole-body robots. Compounding this difficulty, whole-body control often
 207 involves sparse reward signals, which exacerbate the challenges of temporal credit assignment and
 208 policy optimization. Furthermore, prevailing approaches are typically confined to single-task or
 209 single-modality domains in embodied control, raising concerns about their versatility and scalability.

210 4.2 HVD: HIERARCHICAL VALUE-DECOMPOSED OFFLINE RL

212 To address the above issues, we introduce **Hierarchical Value-Decomposed Offline Reinforcement**
 213 **Learning (HVD)**, designed for high-dimensional, whole-body control in robotic systems. Unlike
 214 conventional approaches that decompose the policy (Sentis & Khatib, 2006), HVD introduces
 215 hierarchy directly into the Q-value function through spatial decomposition. This key design allows
 us to maintain a unified policy network while enabling fine-grained, component-specific value

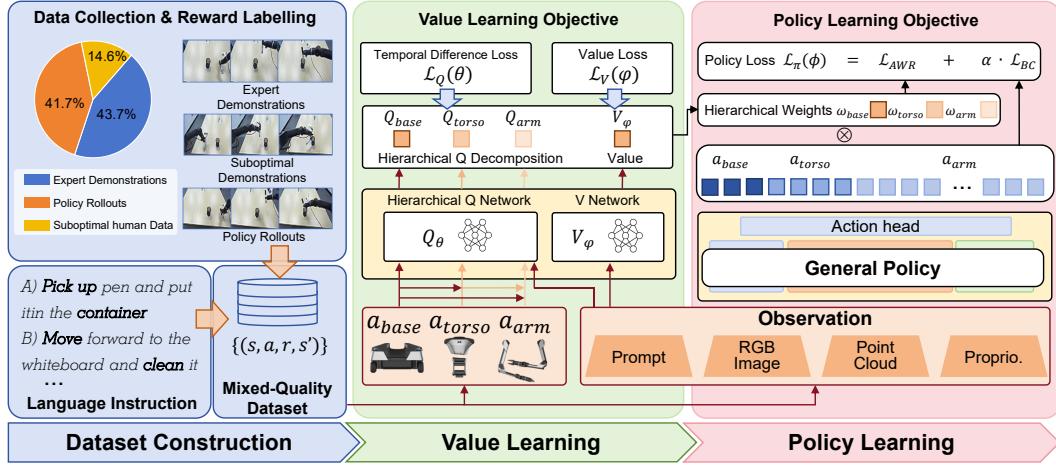


Figure 2: **Learning framework of HVD.** The proposed HVD framework consists of three stages: (1) Dataset construction and reward labeling from expert, suboptimal, and rollout data; (2) Hierarchical value learning with kinematically decomposed Q-functions and temporal chunking; and (3) Policy learning via hierarchical advantage-weighted regression. The diagram illustrates the data flow from multi-modal observations through hierarchical value function learning, and ultimately to policy training.

assessment for different parts of the robot. As a result, HVD achieves more accurate credit assignment across long-horizon, multi-step behaviors for whole-body control.

Q-value Decomposition. Inspired by Smith et al. (2012); Pan et al. (2024); Jiang et al. (2025), the action space of whole-body control policies can be decomposed into three hierarchical components corresponding to its physical structure, $\mathcal{A} = \mathcal{A}_{\text{base}} \times \mathcal{A}_{\text{torso}} \times \mathcal{A}_{\text{arm}}$. At each time step h , the action chunk of size k is defined as $a^{h:h+k} = (a_{\text{base}}^{h:h+k}, a_{\text{torso}}^{h:h+k}, a_{\text{arm}}^{h:h+k})$, where each component represents a sub-action controlling a specific subset of the robot’s degrees of freedom. Specifically, $a_{\text{base}}^{h:h+k}$ governs the lower-body motion (e.g., locomotion or base movement), $a_{\text{torso}}^{h:h+k}$ controls the upper body or torso orientation, and $a_{\text{arm}}^{h:h+k}$ manages the arm movements.

Based on this decomposition, we define hierarchical Q-values over temporal chunks of length k , where each level accumulates value estimates conditioned on progressively more complete subsets of the robot’s action space. Specifically, for a chunk starting at timestep h , we compute:

$$\begin{cases} Q_{\text{base}}^{h:h+k} = Q_{\theta}(s^h, a_{\text{base}}^{h:h+k}), \\ Q_{\text{torso}}^{h:h+k} = Q_{\theta}(s^h, a_{\text{base}}^{h:h+k}, a_{\text{torso}}^{h:h+k}), \\ Q_{\text{arm}}^{h:h+k} = Q_{\theta}(s^h, a_{\text{base}}^{h:h+k}, a_{\text{torso}}^{h:h+k}, a_{\text{arm}}^{h:h+k}). \end{cases} \quad (3)$$

Here, each Q-value corresponds to a specific part of the robot, creating a layered structure for the value function, enabling more precise, joint-level credit assignment.

Hierarchical Value Estimation. To train the hierarchical Q-function in Equation 3, we employ a multi-level TD learning loss that aligns each partial Q-value with its corresponding estimated return:

$$\mathcal{L}_i^h(\theta) = \mathbb{E} \left[(r(s^h, a^{h:h+k}) + V_{\psi}(s^{h+k+1}) - Q_i^{h:h+k})^2 \right], \quad \text{where } i \in \{\text{base, torso, arm}\}. \quad (4)$$

Here, $r(s^h, a^{h:h+k}) = \sum_{j=h}^{h+k} r(s^j, a^j)$ is the reward for executing the action chunk $a^{h:h+k}$ on the state s^h , which aggregates the per-timestep rewards over the sub-episode from time h to $h+k$. Moreover, $V_{\psi}(s^{h+k+1})$ represents the estimated value of the next state predicted by a value network parameterized by ψ . The Q-learning objective is designed to minimize the temporal difference between the predicted Q-value and the target value, and the total Q loss combines all hierarchical levels defined as below:

$$\mathcal{L}_Q(\theta) = \frac{1}{H} \sum_{h=1}^H [\mathcal{L}_{\text{base}}^h(\theta) + \mathcal{L}_{\text{torso}}^h(\theta) + \mathcal{L}_{\text{arm}}^h(\theta)]. \quad (5)$$

270 The result is a value decomposition that promotes credit assignment across space with temporal Q
 271 chunking, enabling more sample-efficient and coordinated whole-body control.
 272

273 **Implicit Value Learning.** Concurrently, we incorporate implicit value learning to align value
 274 estimates across the hierarchical levels of our framework. By leveraging an in-sample learning
 275 paradigm (Wainwright, 2019; Hansen-Estruch et al., 2023), HVD effectively mitigates the risk of
 276 value overestimation caused by OOD actions, a persistent challenge in Q-learning methods:
 277

$$\mathcal{L}_V(\psi) = \frac{1}{H} \sum_{h=1}^H \mathbb{E} \left[\sum_{i \in \{\text{base, torso, arm}\}} \left[\alpha \exp(Q_i^{h:h+k} - V_\psi(s^h)) - \alpha (Q_i^{h:h+k} - V_\psi(s^h)) \right] \right], \quad (6)$$

280 where $\alpha > 0$ is a temperature parameter controlling the strength of the constraint.
 281

282 By optimizing this loss, we establish a soft lower bound on the value estimates across all hierarchical
 283 Q-heads, ensuring that limb-level value predictions remain aligned with global, whole-body goals.
 284

285 **Policy Learning** We train the policy network π_ϕ using a hierarchical variant of Advantage-Weighted
 286 Regression (AWR) (Peters & Schaal, 2007; Peng et al., 2019; Nair et al., 2020). Rather than uniformly
 287 imitating all actions in the dataset, our method assigns importance weights to action chunks based on
 288 estimated advantages, encouraging the policy to prefer high-value behaviors while down-weighting
 289 low-return trajectories. This weighting is applied separately at each hierarchical level, allowing
 290 critical sub-actions to be emphasized even when other components generate lower returns.
 291

292 The per-level advantage weight for an action chunk $a_i^{h:h+k}$ executed from state s^h is defined as:
 293

$$\omega_i^{h:h+k}(s^h, a_i^{h:h+k}) = \frac{\alpha |\exp(\alpha(Q_i^{h:h+k} - V_\psi(s^h))) - 1|}{|Q_i^{h:h+k} - V_\psi(s^h)|}, \quad (7)$$

294 where $\alpha > 0$ controls the sharpness of advantage-based reweighting.
 295

296 This formulation ensures that actions with higher relative advantage receive exponentially increasing
 297 weight, while preserving gradient flow even near the decision boundary. Furthermore, to learn robust
 298 policies from limited demonstrations, we combine two loss terms: i) an RL term trained on an offline
 299 dataset \mathcal{D}^O , weighted by the advantage scores; and ii) a BC term trained on a smaller set of expert
 300 trajectories \mathcal{D}^E , providing a stabilizing prior:
 301

$$\mathcal{L}_\pi^{\text{RL}}(\phi) = \frac{1}{H} \sum_{h=1}^H \sum_i \mathbb{E}_{\mathcal{D}^O} \left[\omega_i^{h:h+k}(s^h, a_i^{h:h+k}) \left\| \epsilon - \pi_\phi(\sqrt{\hat{\alpha}} a_i^{h:h+k} + \sqrt{1 - \hat{\alpha}} \epsilon, s^h, t) \right\| \right], \quad (8)$$

$$\mathcal{L}_\pi^{\text{BC}}(\phi) = \frac{1}{H} \sum_{h=1}^H \sum_i \mathbb{E}_{\mathcal{D}^E} \left[\left\| \epsilon - \pi_\phi(\sqrt{\hat{\alpha}} a_i^{h:h+k} + \sqrt{1 - \hat{\alpha}} \epsilon, s^h, t) \right\| \right], \quad (9)$$

$$\mathcal{L}_\pi(\phi) = \mathcal{L}_\pi^{\text{RL}}(\phi) + \beta \mathcal{L}_\pi^{\text{BC}}(\phi). \quad (10)$$

309 where $\epsilon \sim \mathcal{N}_{\text{pol}}(0, I)$ denotes Gaussian noise, t is the noising timestep, $\hat{\alpha}_t$ represents the noise
 310 schedule parameter in diffusion training, and $\beta > 0$ controls the trade-off between reinforcement
 311 learning-driven exploration and expert imitation.
 312

313 4.3 GENERAL ALGORITHMS AND PRACTICAL IMPLEMENTATION

314 Algorithm 1 offers an overview of the HVD approach, which operates in two phases. The first phase
 315 focuses on hierarchical value learning, where both the value network V_ψ and Q-network Q_θ are
 316 updated using the hierarchical value-decomposed learning loss (Equation 6) and TD loss (Equation
 317 5), respectively. The second phase performs policy extraction, where the policy network π_ϕ is trained
 318 to maximize the cumulative returns through advantage weighted regression (Equation 10).
 319

320 **Model Architecture.** As illustrated in Figure 21, our hierarchical Q-network adopts a unified
 321 multi-modal architecture centered around a Transformer-based backbone. The model can optionally
 322 process a rich set of sensory modalities as input by processing them into token embeddings, including
 323 egocentric RGB images, point cloud data from depth sensors, natural language task instructions, and
 324 proprioceptive state. More detailed implementation can be found in Appendix C.3.

324

5 EXPERIMENTS

326 Our experiments aim to address three core questions. **Q1**: Does HVD consistently surpass imitation
 327 learning baselines across diverse policy architectures? (Section 5.2) **Q2**: Does hierarchical value
 328 decomposition yield more accurate credit assignment, and does this improve the policy performance?
 329 (Section 5.3) **Q3**: Can HVD effectively scale to multi-task settings, leveraging shared value structure
 330 to improve overall performance? (Section 5.4)



344 Figure 3: Illustration of evaluated tasks: (a) *Pen Insert*, (b) *Cup Upright*, (c) *Wipe Board*, (d) *Basket*
 345 *Carry*, (e) *Trash Dispose*.

347

5.1 EXPERIMENT SETTINGS

349 **Robot Platform.** We conduct all experiments on the Galaxeia R1, a real-world wheeled humanoid
 350 robot with a 21-DoF whole-body system. Task demonstrations are collected using JoyLo, a teleopera-
 351 tion interface developed by (Jiang et al., 2025). Guidelines are provided to constrain operators to
 352 generate demonstrations that are easier for the robot to learn. Nevertheless, operator skill levels vary,
 353 resulting in a substantial number of suboptimal demonstrations during data collection.

354 **Task Design.** We design a suite of five representative office tidying tasks (see Figure 3 with details
 355 in Appendix B). The tasks require navigation, dexterous manipulation, and bimanual coordination,
 356 with durations from 40-second single-arm actions to **over 120-second** multi-step sequences involving
 357 coordinated locomotion and dual-arm cooperation. Additionally, we assess task difficulty from
 358 temporal complexity, kinematics, control, and coordination complexity as detailed in Appendix B.6.
 359 This diversity enables rigorous testing of both precision and long-horizon whole-body control.

360 **Baselines.** We develop our HVD framework based on three baselines with different input modal-
 361 ities, including the state-of-the-art VLA model π_0 (Black et al., 2024), the 3D-input model WB-
 362 VIMA (Jiang et al., 2025), and the Diffusion Policy (Chi et al., 2023). We evaluate the performance
 363 of policies trained using the original methods on expert datasets and compare them with policies
 364 trained using HVD on mixed-quality datasets.

366 **Evaluation Metrics.** To enable fine-grained assessment of policy performance, each task is de-
 367 composed into distinct logical stages. We report two primary metrics: *success rate* for task-level
 368 evaluation and *normalized stage score* for stage-level analysis. **Moreover, we introduce perturba-
 369 tions to the task environment background, initial task region, and robot’s initial pose to further challenge
 370 robustness and evaluate the model’s ability to generalize under diverse and realistic variations.** Each
 371 policy was evaluated over 50 independent rollouts per task. Reported success rates and stage scores
 372 are averaged over these rollouts to ensure statistical consistency and fair comparison across methods.

373

5.2 BENCHMARK RESULTS

375 We present the main experimental results of our study, evaluating each method under its optimal
 376 training regime to assess peak performance. HVD is trained on the full mixed-quality dataset, while
 377 imitation learning baselines are trained on the expert-only subset, consistent with their reliance on
 high-quality demonstrations.

As shown in Table 1, our proposed **HVD** consistently outperforms standard imitation learning across all five tasks, yielding higher average success rates across policies. The gains are especially pronounced in challenging manipulation tasks such as *Wipe Board* and *Basket Carry*, where robustness to imperfect initial states and partial observability is essential. Moreover, Figure 4 shows that HVD’s benefits extend beyond task-level success: it achieves higher normalized stage scores on nearly all subtasks. This demonstrates that HVD not only improves final outcomes but also enhances policy reliability throughout the entire execution trajectory.

Method (IL/HVD)	Tasks					Avg SR
	Pen Insert	Cup Upright	Wipe Board	Basket Carry	Trash Dispose	
π_0	0.64/0.86	0.82/0.90	0.18/0.32	0.26/0.44	0.28/0.36	0.44/ 0.58
WB-VIMA	0.52/0.78	0.58/0.82	0.12/0.26	0.10/0.10	0.20/0.32	0.30/ 0.46
DP	0.54/0.64	0.66/0.72	0.00/0.00	0.00/0.08	0.08/0.16	0.26/ 0.32

Table 1: Task-level success rate (SR) of IL and HVD across baseline methods on 5 tasks.

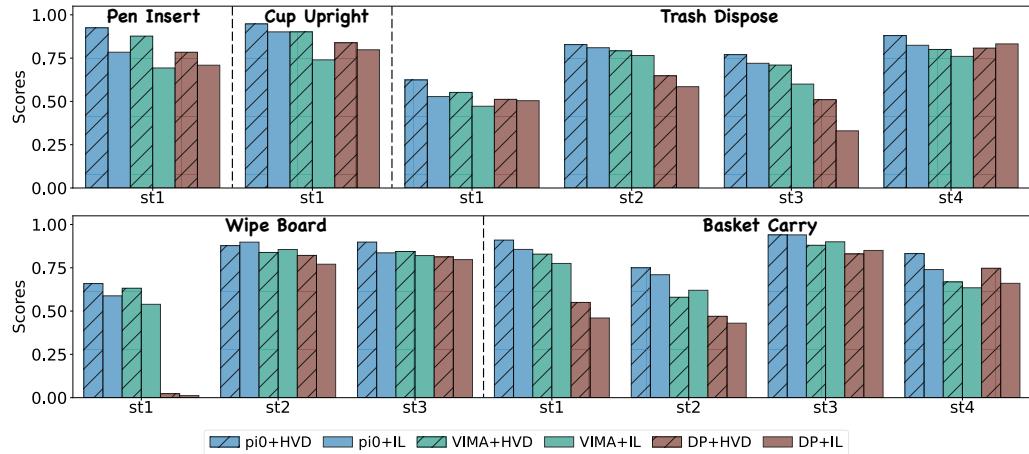


Figure 4: Stage-level scores of IL and HVD across baseline methods across 5 tasks.

5.3 VALUE DECOMPOSITION ABLATION

First, we investigate whether the observed performance gain is primarily attributable to the hierarchical value decomposition or merely to the application of offline RL. To this end, we conduct an ablation study comparing our model with the shared Q-value (named as **HVD w/o Hierarchy**). Both models are trained on the same mixed-quality dataset with the same hyperparameters, ensuring that the only architectural difference is the presence of hierarchical value decomposition. Table 2 shows that removing hierarchical decomposition consistently harms performance across tasks, confirming that the improvements not only stem from the training paradigm, but also from the proposed structure.

Method w/o hierarchy	Tasks					Avg Diff
	Pen Insert	Cup Upright	Wipe Board	Basket Carry	Trash Dispose	
DP	-0.02	0.00	0.00	-0.08	-0.06	-0.03
WB-VIMA	-0.02	0.00	-0.12	-0.08	-0.12	-0.07
π_0	+0.04	-0.02	-0.14	-0.10	-0.04	-0.05

Table 2: Ablation study results on hierarchical value decomposition across 5 tasks. The value here indicates the success rate changes when removing the hierarchy.

To further analyze how decomposition impacts credit assignment, we visualize the advantage weights ω_i of several key frames in the *Basket Carry* task (Figure 5). At the second key frame, where the robot prepares to stand and hold the basket, HVD assigns higher weights to the arm and torso, reflecting the importance of these components. In contrast, HVD w/o Hierarchy produces uniformly high weights across all frames, failing to differentiate subcomponents. These findings indicate that HVD enables more precise credit assignment, which directly contributes to more accurate and reliable action generation. The more visualization examples are shown in Appendix D.1.

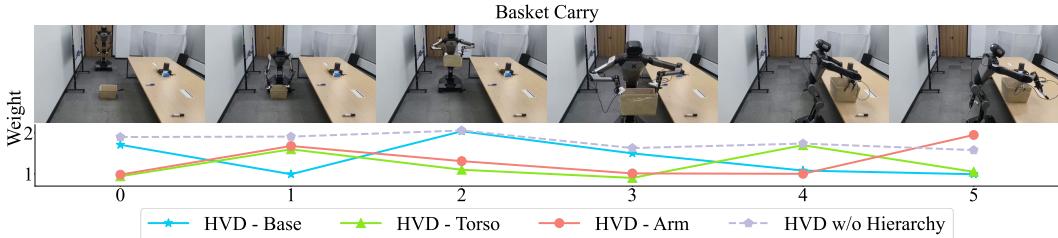


Figure 5: **Credit Assignment Comparison** between HVD and HVD w/o hierarchy.

5.4 MULTI-TASK LEARNING

We evaluate HVD’s ability to scale to multi-task learning by training a single policy on data from all five tasks and comparing its success rate against specialist (single-task) policies. As shown in Table 3, standard IL suffers from negative transfer in the multi-task setting, leading to degraded performance on most tasks. In contrast, HVD mitigates such interference and even surpasses single-task specialists on several tasks. We observe that these gains primarily stem from *enhanced torso robustness* and *more generalizable grasping behaviors* acquired during multi-task training. Together, these results show that HVD effectively leverages shared knowledge across tasks while maintaining specialization, making it more scalable to diverse multi-task settings.

Method	Tasks					Avg SR
	π_0	Pen Insert	Cup Upright	Wipe Board	Basket Carry	
expertise IL		0.64	0.82	0.18	0.26	0.28
multi-task IL		0.50	0.60	0.18	0.24	0.30
expertise HVD		0.86	0.90	0.32	0.44	0.36 (-0.08)
multi-task HVD		0.92	0.94	0.32	0.50	0.60 (+0.02)

Table 3: Task-level success rate comparison of multi-task and expertise policies across 5 tasks.

5.5 COMPARISON WITH OFFLINE RL BASELINES

To further benchmark HVD against other offline RL methods, we implement QIPO (Zhang et al.) using the same base policy π_0 and train in a multi-task learning setting. For a fair comparison, QIPO and HVD share identical network architectures, optimizer configurations, and training hyperparameters. The only differences lie in: (i) the policy update mechanism, and (ii) the use of hierarchical value decomposition. Specifically, QIPO initializes its policy and Q-networks using BC and TD losses, respectively, followed by iterative policy improvement steps. We adopt the hyperparameter recommendations from the original QIPO paper as our starting point and perform minimal tuning to ensure stable convergence. The final configuration uses $\beta = 1$, $M = 16$ (number of sampled actions), and $K_{\text{renew}} = 10$ (policy renewal interval).

As shown in Table 4, HVD consistently outperforms QIPO across all five tasks in terms of success rate (SR). The improvement is most significant in long-horizon, whole-body control tasks, such as *Basket Carry* and *Trash Dispose*, where accurate credit assignment over extended action sequences is critical. We attribute this advantage to HVD’s hierarchical value decomposition, which enables more precise reward propagation and subgoal-aware policy learning.

486 487 488 489 490 491	Method π_0 MultiTask	Tasks					Avg SR 0.51 0.60
		Pen Insert	Cup Upright	Wipe Board	Basket Carry	Trash Dispose	
QIPO	0.90	0.90	0.30	0.20	0.25	0.51	
HVD	0.92	0.94	0.32	0.50	0.30	0.60	

Table 4: Task-level success rate of HVD and QIPO across 5 tasks.

6 RELATED WORK

Whole-body Control Policy Learning. Whole-body control is a central challenge in robotics due to the high dimensionality of articulated bodies and the lack of inherent self-stabilization (Hirai et al., 1998; Grizzle et al., 2009). Classical model-based planning emphasizes kinematic feasibility, stability, and reactive regulation (Sentis & Khatib, 2006; Dietrich et al., 2012; Burget et al., 2013; Kaelbling & Lozano-Pérez, 2013; Dai et al., 2014), but struggles with adaptability and scalability in unstructured tasks. Learning-based approaches optimize control policies from data (Siekmann et al., 2021; Li et al., 2021; Dao et al., 2022; Radosavovic et al., 2024; Cheng et al., 2024), enabling dynamic behaviors difficult to engineer manually (Xia et al., 2021; Jiang et al., 2024; Fu et al., 2024; Arm et al., 2024). Recent advances include generative policies for capturing multimodal action distributions (Fu et al., 2023; Jiang et al., 2025), VLA models that ground control in language and perception (Xu et al., 2024; Ding et al., 2025), and hierarchical policy architectures for managing the complexity of humanoid whole-body control (Hansen et al., 2025; Wei et al., 2025; Fu et al., 2025). Despite these advances, most methods still require high-quality demonstrations, limiting scalability to complex real-world tasks.

Offline RL for Embodied Control. Offline RL has made significant strides in embodied control tasks, enabling robots to learn complex behaviors from pre-collected datasets without requiring expert demonstrations (Levine et al., 2020). Previous works have attempted to learn policies from trajectories generated by human failures or during policy evaluation (Kumar et al., 2021; Mandlekar et al., 2022; Sinha et al., 2022; Bhatia et al., 2023; Luo et al., 2023; Zhou et al., 2023; Ma et al., 2024). However, these methods are largely confined to arm-based manipulators, and their effectiveness in high-DoF whole-body control tasks remains unexplored. Recently, there has been an effort to adapt RL algorithms to mobile manipulators (Hu et al., 2023; Pan et al., 2024). However, these approaches often rely on single-modality inputs, limiting their ability to integrate with generalist policies like VLA models (Kim et al., 2024; Black et al., 2024). Furthermore, most of these methods focus on single-task training, raising concerns about their scalability and generalization to multi-task scenarios.

7 CONCLUSION AND LIMITATIONS

Conclusion. In this paper, we present **HVD**, a framework for learning whole-body robotic control from imperfect, real-world demonstrations. By introducing kinematically aware value decomposition within a multi-modal Transformer architecture, HVD enables stable and scalable policy learning in high-dimensional action spaces using suboptimal offline data. Together with the release of **WB-50**, a 50-hour dataset of realistic teleoperation and rollout trajectories, we demonstrate that structured offline RL can effectively leverage partial successes and failures to achieve robust, generalizable control. The information about resource cost is listed in Appendix E.

Limitations. First, HVD relies on human-annotated rewards, which can be costly. Second, we have not yet explored using failed data for pretraining, which could become a valuable paradigm for improving robustness and scalability in open-ended environments. Future work could also investigate leveraging VLM for automated reward labeling to reduce human effort and enhance scalability.

REFERENCES

Philip Arm, Mayank Mittal, Hendrik Kolvenbach, and Marco Hutter. Pedipulate: Enabling manipulation skills using a quadruped robot’s leg. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5717–5723. IEEE, 2024.

540 Jan Beirlant, Goedele Dierckx, Yuri Goegebeur, and Gunther Matthys. Tail index estimation and an
 541 exponential regression model. *Extremes*, 2(2):177–200, 1999.

542

543 Richard Bellman. Dynamic programming. *science*, 153(3731):34–37, 1966.

544

545 Lucas Beyer, Andreas Steiner, André Susano Pinto, Alexander Kolesnikov, Xiao Wang, Daniel
 546 Salz, Maxim Neumann, Ibrahim Alabdulmohsin, Michael Tschannen, Emanuele Bugliarello, et al.
 547 Paligemma: A versatile 3b vlm for transfer. *arXiv preprint arXiv:2407.07726*, 2024.

548

549 Chethan Anand Bhateja, Derek Guo, Dibya Ghosh, Anikait Singh, Manan Tomar, Quan Vuong,
 550 Yevgen Chebotar, Sergey Levine, and Aviral Kumar. Robotic offline rl from internet videos via
 551 value-function pre-training. In *NeurIPS 2023 Foundation Models for Decision Making Workshop*,
 552 2023.

553

554 Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolò Fusai,
 555 Lachy Groom, Karol Hausman, Brian Ichter, et al. π 0: A vision-language-action flow model for
 556 general robot control. *arXiv*, 2410.24164, 2024.

557

558 Kianté Brantley, Wen Sun, and Mikael Henaff. Disagreement-regularized imitation learning. In
 559 *Proceedings of the 8th International Conference on Learning Representations*, 2020.

560

561 Felix Burget, Armin Hornung, and Maren Bennewitz. Whole-body motion planning for manipulation
 562 of articulated objects. In *2013 IEEE International Conference on Robotics and Automation*, pp.
 563 1656–1662. IEEE, 2013.

564

565 Xuxin Cheng, Yandong Ji, Junming Chen, Ruihan Yang, Ge Yang, and Xiaolong Wang. Expressive
 566 whole-body control for humanoid robots. *arXiv preprint arXiv:2402.16796*, 2024.

567

568 Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake,
 569 and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The
 570 International Journal of Robotics Research*, pp. 02783649241273668, 2023.

571

572 Hongkai Dai, Andrés Valenzuela, and Russ Tedrake. Whole-body motion planning with centroidal
 573 dynamics and full kinematics. In *2014 IEEE-RAS International Conference on Humanoid Robots*,
 574 pp. 295–302. IEEE, 2014.

575

576 Jeremy Dao, Kevin Green, Helei Duan, Alan Fern, and Jonathan Hurst. Sim-to-real learning for
 577 bipedal locomotion under unsensed dynamic loads. In *2022 International Conference on Robotics
 578 and Automation (ICRA)*, pp. 10449–10455. IEEE, 2022.

579

580 Alexander Dietrich, Thomas Wimbock, Alin Albu-Schaffer, and Gerd Hirzinger. Reactive whole-
 581 body control: Dynamic mobile manipulation using a large number of actuated degrees of freedom.
 582 *IEEE Robotics & Automation Magazine*, 19(2):20–33, 2012.

583

584 Pengxiang Ding, Jianfei Ma, Xinyang Tong, Binghong Zou, Xinxin Luo, Yiguo Fan, Ting Wang,
 585 Hongchao Lu, Panzhong Mo, Jinjin Liu, et al. Humanoid-vla: Towards universal humanoid control
 586 with visual integration. *arXiv preprint arXiv:2502.14795*, 2025.

587

588 Dylan J Foster, Adam Block, and Dipendra Misra. Is behavior cloning all you need? understanding
 589 horizon in imitation learning. *arXiv*, 2407.15007, 2024.

590

591 Zipeng Fu, Xuxin Cheng, and Deepak Pathak. Deep whole-body control: learning a unified policy
 592 for manipulation and locomotion. In *Conference on Robot Learning*, pp. 138–149. PMLR, 2023.

593

594 Zipeng Fu, Tony Z Zhao, and Chelsea Finn. Mobile aloha: Learning bimanual mobile manipulation
 595 with low-cost whole-body teleoperation. *arXiv preprint arXiv:2401.02117*, 2024.

596

597 Zipeng Fu, Qingqing Zhao, Qi Wu, Gordon Wetzstein, and Chelsea Finn. Humanplus: Humanoid
 598 shadowing and imitation from humans. In *Conference on Robot Learning*, 2025.

599

600 Jessy W Grizzle, Jonathan Hurst, Benjamin Morris, Hae-Won Park, and Koushil Sreenath. Mabel, a
 601 new robotic bipedal walker and runner. In *2009 American Control Conference*, pp. 2030–2036.
 602 IEEE, 2009.

594 Nicklas Hansen, Jyothir S V, Vlad Sobal, Yann LeCun, Xiaolong Wang, and Hao Su. Hierarchical
 595 world models as visual whole-body humanoid controllers. In *International Conference on Learning
 596 Representations*, 2025.

597

598 Philippe Hansen-Estruch, Ilya Kostrikov, Michael Janner, Jakub Grudzien Kuba, and Sergey Levine.
 599 Idql: Implicit q-learning as an actor-critic method with diffusion policies. *arXiv preprint
 600 arXiv:2304.10573*, 2023.

601 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
 602 recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,
 603 pp. 770–778, 2016.

604

605 Kazuo Hirai, Masato Hirose, Yuji Haikawa, and Toru Takenaka. The development of honda humanoid
 606 robot. In *Proceedings. 1998 IEEE international conference on robotics and automation (Cat. No.
 607 98CH36146)*, volume 2, pp. 1321–1326. IEEE, 1998.

608 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
 609 neural information processing systems*, 33:6840–6851, 2020.

610

611 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
 612 Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

613

614 Jiaheng Hu, Peter Stone, and Roberto Martín-Martín. Causal policy gradient for whole-body mobile
 615 manipulation. *arXiv preprint arXiv:2305.04866*, 2023.

616

617 Yucheng Hu, Yanjiang Guo, Pengchao Wang, Xiaoyu Chen, Yen-Jen Wang, Jianke Zhang, Koushil
 618 Sreenath, Chaochao Lu, and Jianyu Chen. Video prediction policy: A generalist robot policy with
 predictive visual representations. *arXiv preprint arXiv:2412.14803*, 2024.

619

620 Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess,
 621 Adnan Esmail, Michael Equi, Chelsea Finn, Niccolò Fusai, et al. $\pi_{0.5}$: a vision-language-action
 622 model with open-world generalization. *arXiv preprint arXiv:2504.16054*, 2025.

623

624 Yunfan Jiang, Ruohan Zhang, Josiah Wong, Chen Wang, Yanjie Ze, Hang Yin, Cem Gokmen, Shuran
 625 Song, Jiajun Wu, and Li Fei-Fei. Behavior robot suite: Streamlining real-world whole-body
 manipulation for everyday household activities. *arXiv preprint arXiv:2503.05652*, 2025.

626

627 Zhenyu Jiang, Yuqi Xie, Jinhan Li, Ye Yuan, Yifeng Zhu, and Yuke Zhu. Harmon: Whole-body motion
 628 generation of humanoid robots from language descriptions. *arXiv preprint arXiv:2410.12773*,
 629 2024.

630

631 Leslie Pack Kaelbling and Tomás Lozano-Pérez. Integrated task and motion planning in belief space.
The International Journal of Robotics Research, 32(9-10):1194–1227, 2013.

632

633 Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre
 634 Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, et al. Scalable deep reinforcement
 635 learning for vision-based robotic manipulation. In *Conference on robot learning*, pp. 651–673.
 PMLR, 2018.

636

637 Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair,
 638 Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source
 639 vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.

640

641 Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The
 International Journal of Robotics Research*, 32(11):1238–1274, 2013.

642

643 Roger Koenker and Kevin F Hallock. Quantile regression. *Journal of economic perspectives*, 15(4):
 644 143–156, 2001.

645

646 Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Bal-
 647 subramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Irena Gao, et al. Wilds: A
 648 benchmark of in-the-wild distribution shifts. In *International conference on machine learning*, pp.
 5637–5664. PMLR, 2021.

648 Vijay Konda and John Tsitsiklis. Actor-critic algorithms. *Advances in neural information processing*
 649 *systems*, 12, 1999.

650

651 Ilya Kostrikov, Ashvin Nair, and Sergey Levine. Offline reinforcement learning with implicit
 652 q-learning. *arXiv preprint arXiv:2110.06169*, 2021.

653

654 Aviral Kumar, Justin Fu, Matthew Soh, George Tucker, and Sergey Levine. Stabilizing off-policy
 655 q-learning via bootstrapping error reduction. *Advances in neural information processing systems*,
 656 32, 2019.

657

658 Aviral Kumar, Anikait Singh, Stephen Tian, Chelsea Finn, and Sergey Levine. A workflow for offline
 659 model-free robotic reinforcement learning. In *5th Annual Conference on Robot Learning*, 2021.

660

661 Aviral Kumar, Anikait Singh, Frederik Ebert, Mitsuhiro Nakamoto, Yanlai Yang, Chelsea Finn, and
 662 Sergey Levine. Pre-training for robots: Offline rl enables learning new tasks from a handful of
 663 trials. *arXiv preprint arXiv:2210.05178*, 2022.

664

665 Sascha Lange, Thomas Gabel, and Martin Riedmiller. Batch reinforcement learning. In *Reinforcement*
 666 *learning: State-of-the-art*, pp. 45–73. Springer, 2012.

667

668 Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial,
 669 review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.

670

671 Chengshu Li, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-
 672 Martín, Chen Wang, Gabrael Levine, Wensi Ai, Benjamin Martinez, Hang Yin, Michael Lingelbach,
 673 Minjune Hwang, Ayano Hiranaka, Sujay Garlanka, Arman Aydin, Sharon Lee, Jiankai Sun, Mona
 674 Anvari, Manasi Sharma, Dhruva Bansal, Samuel Hunter, Kyu-Young Kim, Alan Lou, Caleb R
 675 Matthews, Ivan Villa-Renteria, Jerry Huayang Tang, Claire Tang, Fei Xia, Yunzhu Li, Silvio
 676 Savarese, Hyowon Gweon, C. Karen Liu, Jiajun Wu, and Li Fei-Fei. Behavior-1k: A human-
 677 centered, embodied ai benchmark with 1,000 everyday activities and realistic simulation. *arXiv*
 678 *preprint arXiv:2403.09227*, 2024.

679

680 Qiyang Li, Zhiyuan Zhou, and Sergey Levine. Reinforcement learning with action chunking. *arXiv*
 681 *preprint arXiv:2507.07969*, 2025.

682

683 Zhongyu Li, Xuxin Cheng, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil
 684 Sreenath. Reinforcement learning for robust parameterized locomotion control of bipedal robots.
 685 In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2811–2817.
 686 IEEE, 2021.

687

688 Songming Liu, Lingxuan Wu, Bangguo Li, Hengkai Tan, Huayu Chen, Zhengyi Wang, Ke Xu, Hang
 689 Su, and Jun Zhu. Rdt-1b: a diffusion foundation model for bimanual manipulation. *arXiv preprint*
 690 *arXiv:2410.07864*, 2024.

691

692 Jianlan Luo, Perry Dong, Jeffrey Wu, Aviral Kumar, Xinyang Geng, and Sergey Levine. Action-
 693 quantized offline reinforcement learning for robotic skill learning. In *Conference on Robot*
 694 *Learning*, pp. 1348–1361. PMLR, 2023.

695

696 Chengzhong Ma, Deyu Yang, Tianyu Wu, Zeyang Liu, Houxue Yang, Xingyu Chen, Xuguang
 697 Lan, and Nanning Zheng. Improving offline reinforcement learning with in-sample advantage
 698 regularization for robot manipulation. *IEEE Transactions on Neural Networks and Learning*
 699 *Systems*, 2024.

700

701 Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-
 702 Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. What matters in learning from offline
 703 human demonstrations for robot manipulation. In *Conference on Robot Learning*, pp. 1678–1690.
 704 PMLR, 2022.

705

706 Ashvin Nair, Abhishek Gupta, Murtaza Dalal, and Sergey Levine. Awac: Accelerating online
 707 reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.

708

709 Guoping Pan, Qingwei Ben, Zhecheng Yuan, Guangqi Jiang, Yandong Ji, Jiangmiao Pang, Houde Liu,
 710 and Huazhe Xu. Roboduet: A framework affording mobile-manipulation and crossembodyment.
 711 *arXiv preprint arXiv:2403.17367*, 6, 2024.

702 Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression:
 703 Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*, 2019.
 704

705 Jan Peters and Stefan Schaal. Reinforcement learning by reward-weighted regression for operational
 706 space control. In *Proceedings of the 24th international conference on Machine learning*, pp.
 707 745–750, 2007.

708 Dean Pomerleau. Efficient training of artificial neural networks for autonomous navigation. *Neural
 709 Computation*, 3(1):88–97, 1991.

710

711 Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets
 712 for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision
 713 and pattern recognition*, pp. 652–660, 2017.

714 Ilija Radosavovic, Tete Xiao, Bike Zhang, Trevor Darrell, Jitendra Malik, and Koushil Sreenath.
 715 Real-world humanoid locomotion with reinforcement learning. *Science Robotics*, 9(89):eadi9579,
 716 2024.

717 Nived Rajaraman, Lin F. Yang, Jiantao Jiao, and Kannan Ramchandran. Toward the fundamental
 718 limits of imitation learning. In *Advances in Neural Information Processing Systems 33*, pp.
 719 2914–2924, 2020.

720

721 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical
 722 image segmentation. In *International Conference on Medical image computing and computer-
 723 assisted intervention*, pp. 234–241. Springer, 2015.

724 Stéphane Ross, Geoffrey J. Gordon, and Drew Bagnell. A reduction of imitation learning and
 725 structured prediction to no-regret online learning. In *Proceedings of the 14th International
 726 Conference on Artificial Intelligence and Statistics*, pp. 627–635, 2011.

727

728 Fumihiro Sasaki and Ryota Yamashina. Behavioral cloning from noisy demonstrations. In *International
 729 Conference on Learning Representations*, 2020.

730

731 Luis Sentis and Oussama Khatib. A whole-body control framework for humanoids operating in
 732 human environments. In *Proceedings 2006 IEEE International Conference on Robotics and
 733 Automation, 2006. ICRA 2006.*, pp. 2641–2648. IEEE, 2006.

734

735 Jonah Siekmann, Yesh Godse, Alan Fern, and Jonathan Hurst. Sim-to-real learning of all common
 736 bipedal gaits via periodic reward composition. In *2021 IEEE International Conference on Robotics
 and Automation (ICRA)*, pp. 7309–7315. IEEE, 2021.

737

738 Samarth Sinha, Ajay Mandlekar, and Animesh Garg. S4rl: Surprisingly simple self-supervision for
 739 offline reinforcement learning in robotics. In *Conference on Robot Learning*, pp. 907–917. PMLR,
 2022.

740

741 Christian Smith, Yiannis Karayiannidis, Lazaros Nalpantidis, Xavi Gratal, Peng Qi, Dimos V
 742 Dimarogonas, and Danica Kragic. Dual arm manipulation—a survey. *Robotics and Autonomous
 743 systems*, 60(10):1340–1353, 2012.

744

745 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 746 Poole. Score-based generative modeling through stochastic differential equations. In *International
 747 Conference on Learning Representations*, 2021.

748

749 Galaxea Team. Galaxea diffusion policy. 2025.

750

751 Stephen Tian, Blake Wulfe, Kyle Sargent, Katherine Liu, Sergey Zakharov, Vitor Campagnolo
 752 Guizilini, and Jiajun Wu. View-invariant policy learning via zero-shot novel view synthesis. In
 753 *Conference on Robot Learning*, pp. 1173–1193. PMLR, 2025.

754

755 Stephen Tu, Alexander Robey, Tingnan Zhang, and Nikolai Matni. On the sample complexity of
 756 stability constrained imitation learning. In *Learning for Dynamics and Control Conference*, pp.
 180–191. PMLR, 2022.

Martin J. Wainwright. Variance-reduced q-learning is minimax optimal. *arXiv*, 1906.04697, 2019.

756 Lai Wei, Xuanbin Peng, Ri-Zhao Qiu, Xuxin Cheng, and Xiaolong Wang. Hmc: Learning het-
757 erogeneous meta-control for contact-rich loco-manipulation. In *Robotics: Science and Systems*,
758 2025.

759 Fei Xia, Chengshu Li, Roberto Martín-Martín, Or Litany, Alexander Toshev, and Silvio Savarese.
760 Relmogen: Integrating motion generation in reinforcement learning for mobile manipulation. In
761 *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4583–4590. IEEE,
762 2021.

763 Xinyu Xu, Yizheng Zhang, Yong-Lu Li, Lei Han, and Cewu Lu. Humanvla: Towards vision-language
764 directed object rearrangement by physical humanoid. *Advances in Neural Information Processing
765 Systems*, 37:18633–18659, 2024.

766 Yanjie Ze, Gu Zhang, Kangning Zhang, Chenyuan Hu, Muhan Wang, and Huazhe Xu. 3d diffusion
767 policy: Generalizable visuomotor policy learning via simple 3d representations. *arXiv preprint
768 arXiv:2403.03954*, 2024.

769 Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language
770 image pre-training. In *Proceedings of the IEEE/CVF international conference on computer vision*,
771 pp. 11975–11986, 2023.

772 Shiyuan Zhang, Weitong Zhang, and Quanquan Gu. Energy-weighted flow matching for offline
773 reinforcement learning. In *The Thirteenth International Conference on Learning Representations*.

774 Gaoyue Zhou, Liyiming Ke, Siddhartha Srinivasa, Abhinav Gupta, Aravind Rajeswaran, and Vikash
775 Kumar. Real world offline reinforcement learning with realistic data source. In *2023 IEEE
776 International Conference on Robotics and Automation (ICRA)*, pp. 7176–7183. IEEE, 2023.

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832 A HARDWARE COMPONENTS

Figure 6: Low-cost JoyLo system and Galaxeia R1 robot.

833 As illustrated in Figure 6, the hardware setup comprises a JoyLo system and a Galaxeia R1 robot.
 834 The JoyLo system integrates 3D-printable arm links, low-cost Dynamixel actuators, and commodity
 835 Joy-Con controllers; its control loop runs at 100 Hz while data are recorded at 10 Hz. Functional
 836 buttons on the right Joy-Con are used to start, pause, save, and discard recordings. Logged modalities
 837 include RGB and depth images, point clouds, joint states, odometry, and action commands. The
 838 Galaxeia R1 platform is equipped with a ZED 2i stereo camera, two Intel RealSense D435i cameras,
 839 and two Galaxeia G1 parallel grippers.

840
841 B TASK DEFINITION842
843 B.1 PEN INSERT

844 **Task Description** This task requires the robot to grasp a marker pen lying on a table and insert it
 845 vertically into a fixed pen holder (diameter 8 cm). The core challenge lies in seamlessly executing
 846 the entire sequence from grasping to insertion. For each trial, both the marker pen and pen holder
 847 are randomly positioned on the tabletop within the robot’s operational workspace and field of view.
 848 To further validate generalization, the robot’s starting position and torso pose are also randomized
 849 within a constrained range.

850
851 **Evaluation Rubric** The task is evaluated as a single, continuous stage focusing on the successful
 852 transfer of the pen to the holder.

853
854 **Stage 1: Grasp and Insert Pen into Holder**

- 855 • 0.0 points: The robot fails to grasp the marker pen (left of Figure 7).
- 856 • 0.5 points: The robot successfully grasps the marker pen but fails to place it in the
 857 holder, for instance, by dropping the pen outside the holder due to an insecure hold (middle
 858 of Figure 7).
- 859 • 1.0 points: The robot firmly grasps the marker pen and successfully places it into the
 860 holder (right of Figure 7).

861
862 **Prompt:** pick the pen and put it into the holder
 863

Figure 7: Scoring rubric visualization for *Pen Insert*. Stage 1: Grasp and Insert Pen into Holder

871 B.2 CUP UPRIGHT

873 **Task Description.** This task requires the robot to grasp a horizontally lying plastic cup (diameter
 874 8 cm) and place it in a stable, upright position onto a target plate. The primary challenges involve
 875 dexterous reorientation of the cup during manipulation and ensuring a steady final placement. For
 876 each trial, the cup is randomly placed on the tabletop, and the robot’s starting pose is randomized
 877 within a constrained range to test for policy generalization. The task demands a combination of
 878 precise grasping and controlled, stable placement.

879 **Evaluation Rubric** The task is evaluated as a single, continuous stage that assesses the entire
 880 sequence from grasping to successful upright placement.

882 Stage 1: Grasp and Place Cup Upright

- 883 • 0.0 points: The robot fails to secure the cup with its gripper (left of Figure 8).
- 884 • 0.5 points: The robot grasps the cup but fails to place it upright on the plate, either due
 885 to losing its grip or improper reorientation (middle of Figure 8).
- 886 • 1.0 points: The robot firmly grasps the cup and places it steadily in an upright position
 887 on the plate (right of Figure 8).

Figure 8: Scoring rubric visualization for *Cup Upright*. Stage 1: Grasp and Place Cup Upright

894 **Prompt:** pick the cup and put it onto the coaster

895 B.3 WIPE BOARD

901 **Task Description.** This task is composed of three sequential stages. In the first stage, the robot
 902 navigates its base to a position in front of the whiteboard and grasps an eraser. The second stage
 903 involves wiping the designated markings from the board. In the third and final stage, the robot places
 904 the eraser back into its designated slot. To promote policy generalization, each trial is initialized with
 905 randomized starting poses for the robot and varied positions for the writing on the whiteboard, both
 906 within a predefined area. This task demands precise physical interaction, as successful execution
 907 hinges on delicate force control: excessive pressure may cause the whiteboard to tilt, while insufficient
 908 force will fail to clean the markings completely.

909 Stage 1: Approach Whiteboard and Grasp Eraser

- 910 • 0.0 points: The robot fails to navigate to a position where any eraser is reachable (left
 911 of Figure 9).
- 912 • 0.5 points: The robot successfully navigates to the whiteboard but fails to establish a
 913 stable grasp on an eraser (middle of Figure 9).
- 914 • 1.0 points: The robot successfully navigates to the whiteboard and executes a stable
 915 grasp on an eraser, suitable for the wiping motion (right of Figure 9).

916 Stage 2: Wipe Markings

918
919
920
921922 Figure 9: Scoring rubric visualization for *Wipe Board*. Stage 1: Approach Whiteboard and Grasp
923 Eraser924
925
926
927
928
929
930
931
932

- 0.0 points: The robot fails to make effective contact with the markings, due to dropping the eraser or significant positioning errors (left of Figure 10).
- 0.5 points: The robot partially erases the markings, or causes the whiteboard to tilt due to imprecise force control or positioning error (middle of Figure 10).
- 1.0 points: The robot completely erases the markings while maintaining stable contact with the whiteboard (right of Figure 10).

933
934
935
936937 Figure 10: Scoring rubric visualization for *Wipe Board*. Stage 2: Wipe Markings
938939
940**Stage 3: Return Eraser**941
942
943
944
945
946
947

- 0.0 points: The robot releases the eraser prematurely before reaching the designated slot (left of Figure 11).
- 0.5 points: The robot reaches the slot but fails to place the eraser correctly due to positioning inaccuracies, causing it to be dropped (middle of Figure 11).
- 1.0 points: The robot successfully and stably places the eraser back into its slot (right of Figure 11).

948
949
950
951
952953 Figure 11: Scoring rubric visualization for *Wipe Board*. Stage 3: Return Eraser
954955
956**Prompt:** move to the whiteboard and clean the whiteboard957
958**B.4 BASKET CARRY**959
960
961
962
963
964
965
966
967

Task description. This task consists of four distinct stages. First, the robot navigates to a cuboid basket placed on the ground. In the second stage, it bends down and lifts the basket using coordinated movements of both arms. For the third stage, the robot turns while holding the basket and places it onto a nearby table. In the final stage, the robot utilizes its left and right arms respectively to pick up markers from the table and place them into the basket. To promote policy generalization, each trial is initialized with the robot at a randomized starting pose, the basket at a varied position on the ground, and the markers at randomized locations on the table, all within a predefined area. This task demands long-horizon planning and effective bimanual coordination.

968
969
970
971**Stage 1: Approach Basket**

- 0.0 points: The robot navigates to an incorrect position, rendering the basket unreach-able for lifting (left of Figure 12).
- 0.5 points: The robot navigates to a misaligned position, preventing a symmetric, bimanual grasp required for a stable lift (middle of Figure 12).

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

- 1.0 points: The robot successfully navigates to a centered position directly in front of the basket, enabling a symmetric, bimanual lift (right of Figure 12).



Figure 12: Scoring rubric visualization for *Basket Carry*. Stage 1: Approach Basket

Stage 2: Lift Basket

- 0.0 points: The robot fails to lift the basket off the ground (left of Figure 13).
- 0.5 points: The robot lifts the basket with both hands but fails to keep it level or properly centered with its body (middle of Figure 13).
- 1.0 points: The robot successfully lifts the basket, maintaining a level and centered orientation relative to its body (right of Figure 13).



Figure 13: Scoring rubric visualization for *Basket Carry*. Stage 2: Lift Basket

Stage 3: Place Basket on Table

- 0.0 points: The robot fails to place the basket onto the table surface (left of Figure 14).
- 0.5 points: The basket is placed on the table but is either dropped from a height or left significantly misaligned with the table's edge (middle of Figure 14).
- 1.0 points: The robot smoothly and squarely places the basket onto the table (right of Figure 14).



Figure 14: Scoring rubric visualization for *Basket Carry*. Stage 3: Place Basket on Table

Stage 4: Place Markers in Basket

- 0.0 points: The robot fails to place any of the markers into the basket (left of Figure 15).
- 0.5 points: The robot successfully places the marker from one side into the basket (middle of Figure 15).
- 1.0 points: The robot successfully places both markers into the basket (right of Figure 15).

Prompt: move to the storage box, pick up the storage box and place it on the table, then put the pen on the table into the box

B.5 TRASH DISPOSE

Task Description This task is composed of four sequential stages, designed to evaluate the robot's capability in long-horizon planning and whole-body coordination within a practical cleanup scenario.

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
Figure 15: Scoring rubric visualization for *Basket Carry*. Stage 4: Place Markers in Basket

In the first stage, the robot presses the top of a tabletop trash can to open its spring-loaded lid. In the second stage, it turns its body to the right and bends down to grasp a crumpled paper towel placed on a nearby chair, an action requiring substantial whole-body coordination to maintain balance. For the third stage, the robot turns back to the left, moves its arm above the trash can opening, and releases the paper towel. In the final stage, the robot must close the lid and press it down again to lock it in place.

To promote policy generalization, each trial initializes with the robot in a randomized starting pose and the paper towel at a varied position on the chair, both within predefined areas. This task presents a significant challenge due to its extended, multi-stage nature. It demands seamless transitions between pressing, grasping, placing, and locking sub-tasks, all while executing complex, coordinated movements.

Stage 1: Open Trash Can Lid

- 0.0 points: The robot fails to open the trash can lid (left of Figure 16).
- 0.5 points: The robot opens the lid but does not succeed on the first attempt (middle of Figure 16).
- 1.0 points: The robot successfully opens the lid on the first attempt (right of Figure 16).

1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
Figure 16: Scoring rubric visualization for *Trash Dispose*. Stage 1: Open Trash Can Lid

Stage 2: Grasp Paper Towel

- 0.0 points: The robot fails to grasp the paper towel from the chair (left of Figure 17).
- 0.5 points: The robot grasps the paper towel but not on the first attempt (middle of Figure 17).
- 1.0 points: The robot successfully grasps the paper towel on the first attempt (right of Figure 17).

1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
Figure 17: Scoring rubric visualization for *Trash Dispose*. Stage 2: Grasp Paper Towel

Stage 3: Dispose of Paper Towel

- 0.0 points: The robot fails to dispose of the paper towel into the trash can (left of Figure 18).
- 0.5 points: The paper towel lands on the rim or gets stuck at the edge of the trash can during disposal (middle of Figure 18).
- 1.0 points: The robot successfully disposes of the paper towel into the trash can (right of Figure 18).

Figure 18: Scoring rubric visualization for *Trash Dispose*. Stage 3: Dispose of Paper Towel

Stage 4: Close and Lock Trash Can Lid

- 0.0 points: The robot fails to close the lid of the trash can (left of Figure 19).
- 0.5 points: The robot pushes the lid down but fails to press it again to lock it, causing the lid to remain unlatched (middle of Figure 19).
- 1.0 points: The robot successfully closes the lid and presses it to ensure it is securely locked (right of Figure 19).

Figure 19: Scoring rubric visualization for *Trash Dispose*. Stage 4: Close and Lock Trash Can Lid

Prompt: Open the trash bin, turn around, pick up the trash on the chair and put it into the bin, then close the bin.

B.6 TASK COMPLEXITY METRIC

To provide a quantitative and objective measure of difficulty for our task suite, we compute a composite complexity score for each task from the collected expert demonstration data ($\mathcal{D}_{\text{expert}}$). This score is derived from four distinct metrics, each capturing a different aspect of task complexity. Let a single expert trajectory be a sequence of 21-DoF joint states $\{q_t\}_{t=1}^T$. The metrics are calculated for each trajectory and then averaged across all demonstrations for a given task.

Temporal Complexity (C_{time}) This metric captures the temporal length of the task and is calculated as the average duration in seconds over all expert demonstrations.

$$C_{\text{time}} = \frac{1}{N} \sum_{i=1}^N T_i \cdot \Delta t, \quad (11)$$

where N is the number of expert trajectories and Δt is the time step duration.

Kinematic Complexity ($C_{\text{kinematic}}$) This metric quantifies the total magnitude of motion required, which is distinct from temporal duration. It is calculated as the average sum of the ℓ_1 norm of joint displacements between consecutive timesteps, capturing the spatial extent of the behavior.

$$C_{\text{kinematic}} = \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^{T_i-1} \|q_{t+1}^{(i)} - q_t^{(i)}\|_1, \quad (12)$$

where $q_t^{(i)}$ is the joint state vector for the i -th trajectory at time t .

Control Complexity (C_{control}) This metric serves as a proxy for control difficulty by measuring the lack of smoothness in the motion. We approximate the average total jerk using the third-order finite difference of the joint positions.

$$C_{\text{control}} = \frac{1}{N} \sum_{i=1}^N \sum_{t=2}^{T_i-2} \|(q_{t+2}^{(i)} - 3q_{t+1}^{(i)} + 3q_t^{(i)} - q_{t-1}^{(i)})\|_2, \quad (13)$$

Higher values indicate more frequent changes in acceleration, suggesting a higher demand on the controller.

1134 **Coordination Complexity** (C_{coord}) This metric estimates the number of joints actively involved in
 1135 the task. For each trajectory, we compute the variance σ_j^2 for each of the 21 joints over time. The
 1136 effective dimensionality is the average number of joints whose variance exceeds a small threshold ϵ
 1137 (e.g., 10^{-4}).

$$1138 \quad C_{\text{coord}} = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^{21} \mathbb{I}(\sigma_j^2 > \epsilon). \quad (14)$$

1141 where $\mathbb{I}(\cdot)$ is the indicator function.

1142

1143 **Complexity Score** To synthesize these individual metrics into a single, comparable score for each
 1144 task, we first perform a min-max normalization on each of the four metrics across the entire task suite.
 1145 This procedure scales the values of each metric to the range of $[0, 1]$, ensuring that each component
 1146 contributes equally to the final score regardless of its original units or scale. The final composite
 1147 complexity score, C_{final} , is then computed as the unweighted average of these four normalized scores,
 1148 providing a holistic and unified measure of task difficulty. The results of this analysis are presented
 1149 in Table 5.

1150 Table 5: Quantitative analysis of task complexity. All presented metric scores (C_k) have been
 1151 min-max normalized to the range $[0, 1]$ for direct comparison. The final score is the unweighted
 1152 average of these normalized values, confirming our task suite spans a graduated range of difficulty.

1153

1154 Task	1155 Temporal (C_{time})	1155 Kinematic ($C_{\text{kinematic}}$)	1155 Control (C_{control})	1155 Coordination (C_{coord})	1155 Final Score (C_{final})
1157 Pen insert	0.00	0.00	0.05	0.01	0.02
1158 Cup upright	0.06	0.51	0.00	0.00	0.14
1159 Wipe board	0.63	0.73	1.00	0.06	0.61
1160 Basket carry	0.47	1.00	0.86	0.78	0.78
1161 Trash dispose	1.00	0.92	0.36	1.00	0.82

1162

1163

B.7 DATASET DESCRIPTION

1164

1165

The WB-50 dataset is a reward-labeled whole-body manipulation dataset spanning over 50 hours of diverse robot experience. WB-50 contains three data sources to reflect realistic data distributions: (i) expert demonstrations (43.7%), (ii) imperfect teleoperation (14.6%), and (iii) policy rollouts (41.7%) — the latter two comprising the majority, mirroring the scarcity of perfect supervision in practice. The proportion of the five task data frames to the total number of frames in the dataset is shown in the Figure 20. It also shows the successful and failed trajectories generated during the expert data collection process of each task, and the number of frames of successful and failed trajectories generated by imitation learning strategy reasoning is also shown. Moreover, we assign rewards of 0, 0.5, and 1 at the end of each subtask according to the degree of task completion, and apply a step penalty of -0.001 for all other steps.

1175

1176

C IMPLEMENTATION DETAILS

1177

1178

C.1 ALGORITHM PIPELINE

1179

We present the pseudocode of our method as in Algorithm 1.

1180

1181

C.2 IMITATION LEARNING BASELINES

1182

1183

WB-VIMA. Our implementation of WB-VIMA is based on the official policy codebase (Jiang et al., 2025) and applies minor adjustments to the model and training hyperparameters. The detailed parameters are summarized in Table 6.

1184

1185

1186

1187

Diffusion Policy. We build our diffusion policy on the official GalaxeDP codebase (Team, 2025), which demonstrates strong compatibility with the Galaxe R1 robot, also developed by Galaxe. As

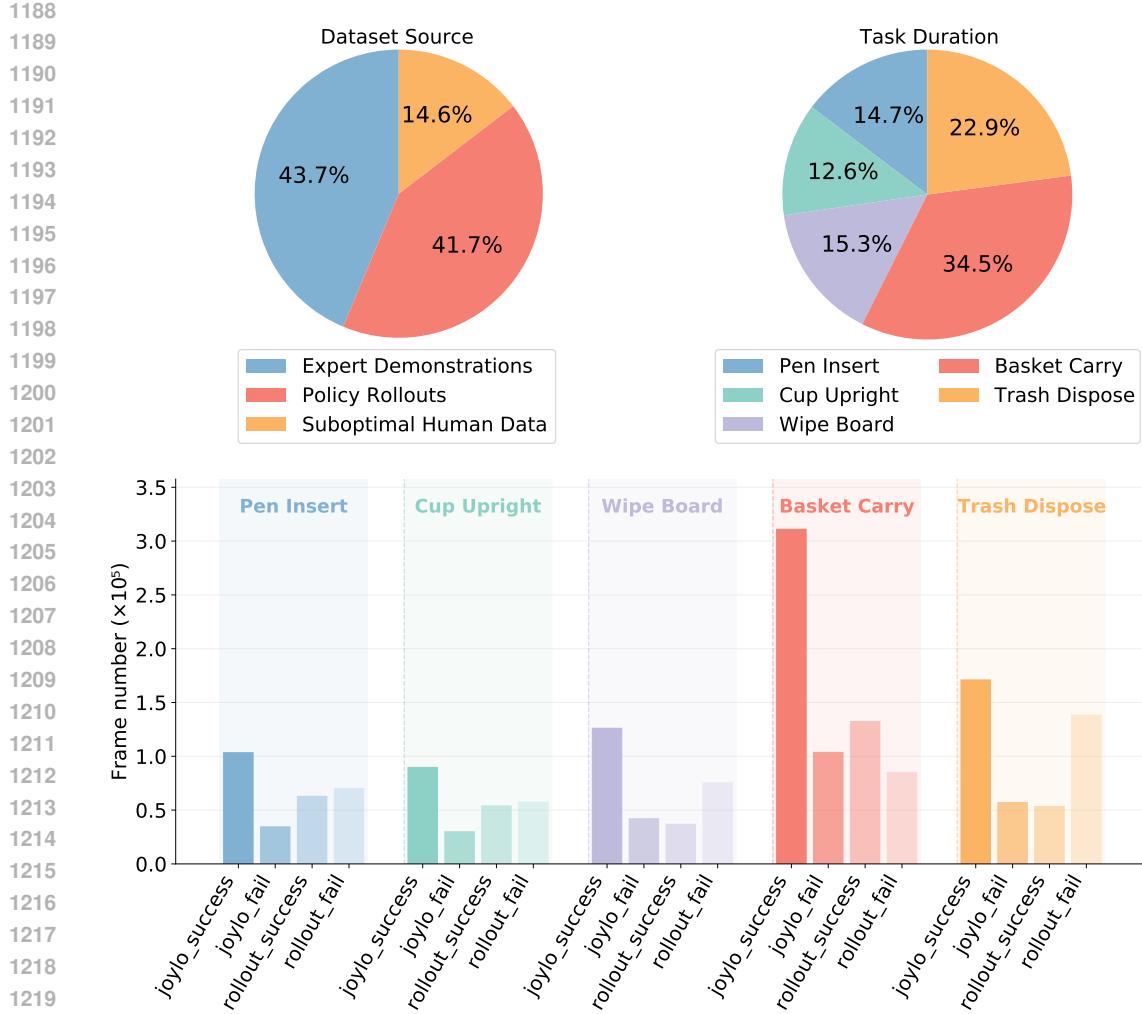


Figure 20: Data distribution of **WB-50**. The plot shows the proportional composition of the dataset by source, including expert demonstrations, suboptimal expert trajectories and rollout data, together with counts of frames of successful and failed episodes collected during data acquisition.

Algorithm 1 Hierarchical Value-Decomposed Offline Reinforcement Learning (HVD)

```

1: Input: Offline dataset  $\mathcal{D}$ , action hierarchy:  $\mathcal{A} = \mathcal{A}_{\text{base}} \times \mathcal{A}_{\text{torso}} \times \mathcal{A}_{\text{arm}}$ 
2: Initialize value network  $V_\psi$ , Q-network  $Q_\theta$ , policy network  $\pi_\phi$ 
   {Phase 1: Hierarchical Value Learning}
3: for each gradient step do
4:   (Update value network)
5:    $\psi \leftarrow \psi - \lambda_V \nabla_\psi \mathcal{L}_V(\psi)$  by Equation equation 6
6:   (Update Q-network)
7:    $\theta \leftarrow \theta - \lambda_Q \nabla_\theta \mathcal{L}_Q(\theta)$  by Equation equation 5
8: end for
   {Phase 2: Policy Extraction}
9: for each gradient step do
10:  (Update policy network)
11:   $\phi \leftarrow \phi - \lambda_\pi \nabla_\phi \mathcal{L}_\pi(\phi)$  by Equation equation 10
12: end for
13: Output: Trained policy  $\pi_\phi$ 

```

1242
1243
1244 Table 6: Hyperparameters of WB-VIMA model.
1245
1246
1247
1248
1249
1250
1251

Hyperparameter	Value	Hyperparameter	Value	Hyperparameter	Value
PointNet		Prop. MLP		Transformer	
N_{pcd}	4096	Input Dim	21	Embed Size	512
Hidden Dim	256	Hidden Dim	256	Num Layers	8
Hidden Depth	2	Hidden Depth	3	Num Heads	8
Output Dim	256	Output Dim	256	Drop Rate	0.1
Activation	GELU	Activation	ReLU	Activation	GEGLU

1252
1253 Table 7: Hyperparameters of WB-VIMA training process.
1254
1255
1256
1257
1258
1259
1260

Hyperparameter	Value
Learning Rate	1×10^{-4}
Weight Decay	0.1
Learning Rate Warm Up Steps	1000
Learning Rate Cosine Decay Steps	300,000
Minimal Learning Rate	5×10^{-6}

1261
1262 the original implementation supports only 14-DoF dual-arm tasks, we extend it by incorporating
1263 additional control for the torso and mobile base to enable 21-DoF whole-body control tasks. The
1264 hyperparameters and model architectures used in our experiments are summarized in Table 8.
1265

Hyperparameter	Value
Batch Size	32
Chunk Size	20
History Size	2
Learning Rate	1e-4
LR Scheduler	cosine
Optimizer	AdamW
AdamW Betas	[0.9, 0.95]
Weight Decay	1e-4
Max Training Steps	100,000
Image Type	RGB
Egocentric Perception Type	Joint
Observation Encoder	ResNet-18 (He et al., 2016)
Diffusion Model	DDPM (Ho et al., 2020)
Diffusion Steps	20
Diffusion Network	U-Net (Ronneberger et al., 2015)
U-Net Structure	[256, 1024, 4096]

1283
1284 Table 8: Hyperparameters of Diffusion Policy.
1285
1286
1287 π_0 . We adopt the official implementation of π_0 (Black et al., 2024) as our codebase. Key hyperparameters are listed in Table 9.1288
1289 C.3 HVD IMPLEMENTATION DETAILS1290
1291 **Model Architecture.** HVD is implemented using a transformer backbone, which naturally accommodates
1292 multiple input modalities. Specifically, the model processes observations through specialized
1293 encoders for each modality:1294
1295

- **RGB Input:** Three egocentric RGB views (front, left, and right) are independently processed using SigLIP (Zhai et al., 2023), producing sequences of visual tokens that capture spatial context and object semantics.

Hyperparameter	Value
Batch Size	32
Chunk Size	20
Learning Rate	1e-4
LR Scheduler	cosine
Optimizer	AdamW
AdamW Betas	[0.9, 0.95]
AdamW Epsilon	1e-8
Weight Decay	1e-10
Max Training Steps	50,000
Fine-tune Method	LoRA (Hu et al., 2022)

Table 9: Hyperparameters of π_0 .

- **Point Cloud Input:** Depth-derived point clouds are encoded with PointNet (Qi et al., 2017) modules, enabling robust perception of 3D geometry and scene layout, which is particularly beneficial for navigation and object manipulation.
- **Task Instruction:** Natural language commands (e.g., “clean the whiteboard”) are tokenized and embedded to provide high-level goal guidance.
- **Proprioception:** Joint angles, velocities, and end-effector poses are concatenated and normalized to form a compact state vector representing the robot’s internal configuration.

The tokens produced by all modalities are concatenated along the sequence dimension to form a unified representation. Attention masks regulate cross-modal interactions, after which the integrated token sequence is processed by a pretrained PaliGemma model (Beyer et al., 2024). The resulting representations are then passed through an MLP-based value decoder to estimate Q-values for the different hierarchical components. The overall model architecture is depicted in Figure 21.

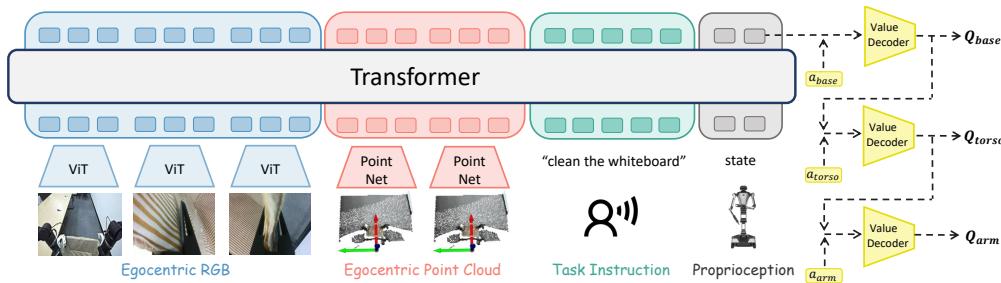


Figure 21: Overall Model Architecture of Hierarchical Q-Network.

Hyperparameters. In our implementation, the parameter α in equation 6 controls the relative weighting between the TD loss and the BC loss, effectively balancing value estimation and policy imitation, while β in equation 10 is used in the exponential weighting of advantages when computing hierarchical action weights, modulating the sensitivity to high-advantage actions. Both α and β are set to their default values of 1.0. Importantly, we did not perform any hyperparameter tuning, yet our method already achieves strong performance, highlighting the effectiveness and robustness of the proposed HVD approach. Configurations and hyperparameter settings are listed in Table 10.

D ADDITIONAL EXPERIMENTS

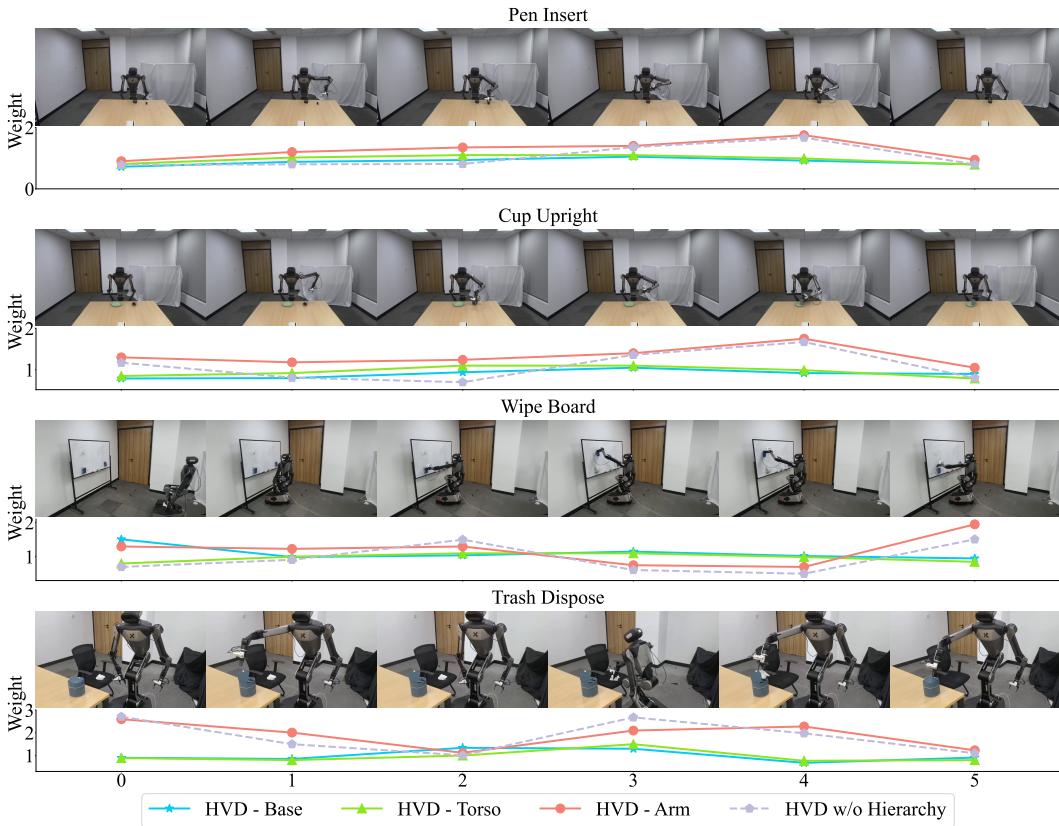
D.1 QUALITATIVE VISUALIZATION OF HIERARCHICAL CREDIT ASSIGNMENT

The remaining weight visualization results are presented in Figure 22. We observe that the HVD w/o Hierarchy weights exhibit a trend similar to the HVD arm weights across all tasks. This indicates that while HVD w/o Hierarchy is able to capture key frames in which the arms are about to move, it

1350	Hyperparameter	Value
1351	Value Network	MLP[128 \times 256 \times 64]
1352	Max Training Steps	30,000
1353	BC Loss Weight α	1
1354	Exponential Weight β	1
1355	Image Encoder	SigLIP (Zhai et al., 2023)
1356	Point Cloud Encoder	PointNet (Qi et al., 2017)
1357	Transformer Model	PaliGemma (Beyer et al., 2024)
1358	Width	256
1359	Depth	4
1360	MLP Dim	1024
1361	Number of Heads	4
1362	Number of KV Heads	1

Table 10: Hyperparameters of HVD.

fails to effectively recognize the contributions of the base and torso. These results further support our conclusion that HVD provides more accurate credit assignment.

Figure 22: **Remaining Credit Assignment Comparison** between HVD and HVD w/o hierarchy.

D.2 QUANTITATIVE ANALYSIS OF HIERARCHICAL CREDIT ASSIGNMENT

To validate the interpretability of our hierarchical framework, we analyzed the component weights during policy training across different task stages. The quantitative results in Table 11 reveal distinct stage-aware credit assignment patterns that align with functional requirements. These results

demonstrate that our hierarchical framework successfully differentiates functional requirements across task stages and allocates appropriate credit to different robot components.

Table 11: Hierarchical component weights across task stages.

Task	Stage	Component Weight		
		Base	Torso	Arm
Wipe Board	Stage 1: Approach and Grasp Eraser	0.962	0.884	1.321
	Stage 2: Wipe Markings	0.946	0.890	1.054
	Stage 3: Return Eraser	0.857	0.826	1.149
Basket Carry	Stage 1: Approach and Grasp Basket	0.980	0.999	1.405
	Stage 2: Lift Basket	0.953	0.909	1.254
	Stage 3: Place Basket on Table	0.837	0.807	1.063
	Stage 4: Place Markers in Basket	0.972	0.980	1.243
Trash Dispose	Stage 1: Open Trash Can Lid	0.933	0.918	1.382
	Stage 2: Grasp Paper Towel	0.930	0.960	1.273
	Stage 3: Dispose of Paper Towel	0.944	0.902	1.356
	Stage 4: Close and Lock Trash Can Lid	0.804	0.813	1.323

D.3 ANALYSIS OF TASK COMPLEXITY AND POLICY PERFORMANCE

By correlating the difficulty metrics presented in Table 5 with the actual success rates of learned policies in Table 1, we can identify several key bottlenecks in current whole-body control approaches and point out the primary sources of performance gains introduced by HVD. Specifically, our analysis reveals that high control complexity and high kinematic coordination demands are the primary failure modes for standard IL, which HVD successfully mitigates.

Overcoming the Smoothness Bottleneck in High-Control Complexity Tasks. C_{control} serves as the most correlated predictor of failure for baseline methods. For instance, the *Wipe Board* task exhibits the maximum Control Complexity ($C_{\text{control}} = 1.00$), indicating a requirement for frequent acceleration changes and high jerk. Standard IL baselines collapse on this task. In contrast, HVD significantly improves π_0 performance with a +77% relative improvement, demonstrating that HVD effectively models the non-smooth, high-frequency dynamics often lost in standard IL training.

Robustness in High-Dimensional Coordination. IL baselines often fail in the tasks characterized by high $C_{\text{kinematic}}$ and C_{coord} due to the difficulty of coordinating high-dimensional joints over large spatial displacements. The *Basket Carry* task, which possesses the highest Kinematic Complexity ($C_{\text{kinematic}} = 1.00$) and very high Coordination Complexity ($C_{\text{coord}} = 0.78$), illustrates this barrier. While the WB-VIMA and DP baselines struggle significantly (with DP failing almost completely), HVD provides its most robust improvement here, lifting the π_0 success rate from 0.26 to 0.44. This indicates that HVD acts as a superior regularizer in high-variance regimes, maintaining structural integrity over wide state spaces where standard IL fails to generalize.

D.4 SIMULATOR EXPERIMENTS

We introduce a simulation experiment on the BEHAVIOR-1K platform to empirically validate our approach, directly addressing the reviewer's query (W5) for a concrete simulated environment.

D.4.1 SIMULATOR PLATFORM

We conduct our empirical evaluation using the BEHAVIOR-1K simulation environment (Li et al., 2024). This platform is a high-fidelity, standardized benchmark focused on humanoid robotics tasks grounded in real-world human needs. The simulated robot model features a base, torso, and arm structure that necessitates complex whole-body coordination, aligning directly with our methodological requirements for addressing the high-Dimensionality of the action space. Figure 23 provides an overview of the simulator setup.



Figure 23: Overview of the simulator in the BEHAVIOR-1K.

Task Design and Decomposition. We selected the Picking Up Trash task from the official BEHAVIOR-1K benchmark as our test environment. This task highly demands robust whole-body control capabilities. Given the computational intensity of long-horizon rollouts in this complex simulator, we constrained each episode to 10,000 control steps. While this limitation prevents full task completion, it is sufficient to cover the most critical phases. To facilitate focused evaluation and learning, we decomposed the task into two distinct, sequential stages with clear success criteria:

- **Approaching:** Successful if the robot positions itself directly in front of the trash bin, such that the bin is visible in at least one camera view.
- **Grasping:** Successful if the robot lifts the trash bin off the ground using either arm.

For reward labeling, successful trajectories (meeting stage criteria) receive a terminal reward of +1.0, provided at the end of the trajectory. Additionally, every intermediate frame incurs a small step penalty of -0.001 to encourage efficient execution. The overall task pipeline, illustrating this two-stage structure, is presented in Figure 24.



Figure 24: Task pipeline for the Picking Up Trash task in BEHAVIOR-1K, illustrating the approaching and grasping stages evaluated in our experiments.

Dataset and Evaluation. Our experimental analysis utilized a training dataset of 300 trajectories: 200 expert demonstrations from the official BEHAVIOR-1K benchmark, supplemented by 100 policy rollouts collected using the pre-trained π_0 -IL baseline. During evaluation, we tested each policy on **100 trials** across 20 unseen scenarios to ensure robustness and generalizability. Table 13 summarizes the performance of the baseline Imitation Learning (π_0 -IL) policy and our Hierarchical Value Decomposition (π_0 -HVD) method across both evaluation stages. HVD consistently demonstrates a significant performance improvement over the baseline:

Implementation Details. The HVD implementation was built directly on top of the official π_0 codebase provided by the BEHAVIOR-1K benchmark, ensuring a fair and consistent comparison with the π_0 , IL baseline. All implementation details, including the hierarchical architecture, network configurations, and hyperparameter settings, remain the same as those introduced in Section C.3, with only minor adjustments as summarized in Table 12.

Task Results. Table 13 summarizes the evaluation outcomes, showing that HVD substantially outperforms the IL baseline. In Stage 1, HVD achieves a **77.0%** success rate than IL's **45.0%**. The advantage becomes even more pronounced in the more challenging grasping stage, where HVD nearly doubles the success rate, reaching **22.0%** compared to IL's **5.0%**. These results highlight HVD's effectiveness in handling complex, high-dimensional whole-body coordination tasks.

Hyperparameter	Value
Max Value Training Steps	10,000
Max Policy Training Steps	80,000
Hierarchical Order	[[0-2], [3-6], [7-22]]

Table 12: Adjusted hyperparameters of HVD in BEHAVIOR-1K.

Method	Stage 1 (SR)	Stage 2 (SR)
π_0 -IL	45.0%	5.0%
π_0 -HVD	77.0%	22.0%

Table 13: Results for Picking Up Trash task in BEHAVIOR-1K.

D.5 ABLATION STUDY ON DATA DISTRIBUTION

To further empirically validate the data distribution robustness of HVD, we conduct an additional ablation on the Pen Insert task under varying expert ratios. During training, we keep the total number of demonstrations fixed at 100 with expert ratios of 20%, 50%, and 80%. During training, π_0 +IL uses only the expert demonstrations, while π_0 +HVD uses all available demonstrations, including both expert and suboptimal data. The results are shown in Table 14, where HVD consistently outperforms IL, confirming that HVD maintains its effectiveness across a wide spectrum of data compositions. Due to time constraints, each policy is evaluated over 20 trials.

Table 14: Success rates on the Pen Insert task under varying expert-to-imitation data ratios.

Method	20 exp + 80 imp	50 exp + 50 imp	80 exp + 20 imp
π_0 +IL	0.15	0.35	0.50
π_0 +HVD	0.55	0.60	0.75

D.6 ABLATION STUDY ON DECOMPOSITION ORDER

To validate the impact of different decomposition orders and validate our design choice, we conduct an ablation experiment. Specifically, we compare two alternative sequencing strategies: (i) Arm–Torso–Base and (ii) Base–Arm–Torso, using the same base policy π_0 across all tasks within a multi-task learning setup. Due to time constraints, the baseline policy is evaluated over 20 trials per task. The results are summarized in the Table 15, clearly demonstrating that our proposed decomposition order outperforms the alternatives, highlighting the importance of the hierarchical structure in our approach.

D.7 COMPARISON WITH RESIDUAL VALUE DECOMPOSITION

To further validate the residual approach versus our proposed independent decomposition, we add an additional baseline across all tasks within a multi-task learning setup, which learns the residual Q-function Q_{residual} using the formulation defined above. Due to time constraints, the baseline policy is evaluated over 20 trials per task. The results are as shown in Table 16, showing that our independent decomposition outperforms the residual variant baseline, achieving a higher average success rate.

E RESOURCE COST

E.1 REWARD LABELING COST

Reward signals are manually annotated according to the task definitions provided in Appendix B. For each demonstration, operators carefully review the recordings from multiple camera views (head, left wrist, and right wrist) and mark the key frames where the robot’s actions satisfy predefined scoring criteria. This process is highly labor-intensive: labeling the entire **WB-50** dataset required

1566
1567
1568 Table 15: Success rates (SR) for different decomposition orders across tasks.
1569
1570
1571

Method	Pen Insert	Cup Upright	Wipe Board	Basket Carry	Trash Dispose	Avg SR
Arm–Torso–Base	0.55	0.60	0.15	0.20	0.10	0.32
Torso–Arm–Base	0.80	0.80	0.20	0.35	0.20	0.47
Base–Torso–Arm (ours)	0.92	0.94	0.32	0.50	0.33	0.60

1572

1573 Table 16: Comparison between residual and independent decomposition variants of HVD.
1574

1575

Method	Pen Insert	Cup Upright	Wipe Board	Basket Carry	Trash Dispose	Avg SR
HVD (residual)	0.75	0.80	0.25	0.50	0.35	0.53
HVD (ours)	0.92	0.94	0.32	0.50	0.33	0.60

1576

1577

1578

1579

1580 two well-trained operators working for nearly 30 hours. The labeled data are subsequently used to
1581 train HVD, yielding improved policy performance.
1582

1583

E.2 COMPUTATIONAL COST

1584

1585

1586

1587

1588

1589

1590

1591

1592

F THE USE OF LARGE LANGUAGE MODELS (LLMs)

1593

1594

1595

1596

1597

1598

1599

1600

1601

1602

1603

1604

1605

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

1618

1619