000 **REINFORCEMENT LEARNING WITH ACTION SEQUENCE** 001 FOR DATA-EFFICIENT ROBOT LEARNING 002 003

Anonymous authors

004

006

008 009

010 011

012

013

014

015

016

017

018

019

021

023

025

026

029

031

037

040

041

042

Paper under double-blind review

ABSTRACT

Training reinforcement learning (RL) agents on robotic tasks typically requires a large number of training samples. This is because training data often consists of noisy trajectories, whether from exploration or human-collected demonstrations, making it difficult to learn value functions that understand the effect of taking each action. On the other hand, recent behavior-cloning (BC) approaches have shown that predicting *a sequence of actions* enables policies to effectively approximate noisy, multi-modal distributions of expert demonstrations. Can we use a similar idea for improving RL on robotic tasks? In this paper, we introduce a novel RL algorithm that learns a critic network that outputs *Q*-values over a sequence of actions. By explicitly training the value functions to learn the consequence of executing a series of current and future actions, our algorithm allows for learning useful value functions from noisy trajectories. We study our algorithm across various setups with sparse and dense rewards, and with or without demonstrations, spanning mobile bi-manual manipulation, whole-body control, and tabletop manipulation tasks from BiGym, HumanoidBench, and RLBench. We find that, by learning the critic network with action sequences, our algorithm outperforms various RL and BC baselines, in particular on challenging humanoid control tasks.



Figure 1: Summary of results. Coarse-to-fine Q-Network with Action Sequence (CQN-AS) is a value-based RL algorithm that learns a critic network with action sequence. We study CQN-AS on 53 robotic tasks from BiGym (Chernyadev et al., 2024), HumanoidBench (Sferrazza et al., 2024), and RLBench (James et al., 2020), where prior model-free RL algorithms struggle to achieve competitive performance. We show that CQN-AS outperforms various RL and BC baselines such as CQN (Seo et al., 2024), DrQ-v2+ (Yarats et al., 2022), SAC (Haarnoja et al., 2018), and ACT (Zhao et al., 2023).

043 044

045 046

1 INTRODUCTION

047 Reinforcement learning (RL) holds the promise of continually improving policies through online trial-048 and-error experiences (Sutton & Barto, 2018), making it an ideal choice for developing robots that 049 can adapt to various environments. However, despite this promise, training RL agents on robotic tasks 050 typically requires a prohibitively large number of training samples (Kalashnikov et al., 2018; Herzog 051 et al., 2023), which becomes problematic as deploying robots often incurs a huge cost. Therefore many of the recent successful approaches on robot learning have been based on behavior-cloning 052 (BC; Pomerleau 1988), which can learn strong policies from offline expert demonstrations (Brohan et al., 2023b;a; Zhao et al., 2023; Chi et al., 2023; Team et al., 2024; Fu et al., 2024a).



Figure 2: **Coarse-to-fine Q-network with action sequence.** (Left) Our key idea is to train a critic network to output Q-values over *a sequence of actions*. We design our architecture to first obtain features for each sequence step and aggregate features from multiple sequence steps with a recurrent network. We then project these outputs into Q-values at level *l*. (Right) For action inference, we repeat the procedure of computing Q-values for level $l \in \{1, ..., L\}$. We then find the action sequence with the highest Q-values from the last level *L*, and use it for controlling robots at each time step.

One cause for the poor data-efficiency of RL algorithms on robotic tasks is that training data consists 072 of noisy trajectories. When collecting data for training RL agents, we typically inject some noise into 073 actions for exploration (Sehnke et al., 2010; Lillicrap et al., 2016) that may induce trajectories with 074 jerky motions. Moreover, we often initialize training with human-collected demonstrations that can 075 consist of noisy multi-modal trajectories (Chernyadev et al., 2024). Such noisy data distributions 076 make it difficult to learn value functions that should understand the consequence of taking each 077 action. We indeed find that prior RL algorithms perform much worse than the BC baseline on mobile 078 bi-manual manipulation tasks with *human-collected* demonstrations when compared to a tabletop 079 manipulation setup with synthetic demonstrations collected via motion-planning (see Figure 1).

On the other hand, recent BC approaches have shown that *predicting a sequence of actions* enables 081 policies to effectively approximate the noisy, multi-modal distribution of expert demonstrations (Zhao 082 et al., 2023; Chi et al., 2023). Inspired by this, in this paper, we investigate how to use a similar 083 idea for improving the data-efficiency of RL algorithms on robotic tasks. In particular, we present 084 a novel RL algorithm that learns a critic network that outputs Q-values over a sequence of actions 085 (see Figure 2). By training the critic network to explicitly learn the consequence of taking a series of current and future actions, our algorithm enables the RL agents to effectively learn useful value functions from noisy trajectories. We build this algorithm upon a recent value-based RL algorithm 087 that learns RL agents to zoom-into continuous action space in a coarse-to-fine manner (Seo et al., 880 2024), thus we refer to our algorithm as Coarse-to-fine Q-Network with Action Sequence (CQN-AS). 089

090 To evaluate the generality and capabilities of CQN-AS, we study CQN-AS on various setups with 091 sparse and dense rewards, and with or without demonstrations. In BiGym benchmark (Chernyadev 092 et al., 2024), which provides human-collected demonstrations for mobile bi-manual manipulation tasks, CQN-AS outperforms various model-free RL and BC baselines (Yarats et al., 2022; Zhao 093 et al., 2023; Seo et al., 2024). Moreover, in HumanoidBench (Sferrazza et al., 2024), which consists 094 of densely-rewarded humanoid control tasks, we show that CQN-AS can also be effective without 095 demonstrations, outperforming prior model-free RL baselines (Haarnoja et al., 2018; Seo et al., 2024). 096 Finally, in RLBench (James et al., 2020), which provides synthetic demonstrations generated via motion-planning, CQN-AS achieves similar performance as model-free RL and BC baselines on 098 most tasks, but significantly better performance on several long-horizon manipulation tasks. 099

100

071

2 BACKGROUND

101 102

Problem setup We mainly consider a robotic control problem which we formulate as a partially observable Markov decision process (Kaelbling et al., 1998; Sutton & Barto, 2018). At each time step t, an RL agent encounters an observation o_t , executes an action a_t , receives a reward r_{t+1} , and encounters a new observation o_{t+1} from the environment. Because the observation o_t does not contain full information about the internal state of the environment, in this work, we use a stack of past observations as inputs to the RL agent by following the common practice in Mnih et al. (2015). For simplicity, we omit the notation for these stacked observations. When the environment is fully observable, we simply use o_t as inputs. Our goal in this work is to train a policy π that maximizes the expected sum of rewards through RL while using as few samples as possible, optionally with access to a modest amount of expert demonstrations collected either by motion-planners or by humans.

112

130

132 133

138

139 140

141

142

143 144

145

146 147

148 149

Inputs and encoding Given visual observations $\mathbf{o}_t^v = \{\mathbf{o}_t^{v_1}, ..., \mathbf{o}_t^{v_M}\}$ from M cameras, we encode each $\mathbf{o}_t^{v_i}$ using convolutional neural networks (CNN) into $\mathbf{h}_t^{v_i}$. We then process them through a series of linear layers to fuse them into \mathbf{h}_t^v . If low-dimensional observations $\mathbf{o}_t^{\mathbf{1}\mathbf{o}\mathbf{w}}$ are available along with visual observations, we process them through a series of linear layers to obtain $\mathbf{h}_t^{\mathbf{1}\mathbf{o}\mathbf{w}}$. We then use concatenated features $\mathbf{h}_t = [\mathbf{h}_t^v, \mathbf{h}_t^{\mathbf{1}\mathbf{o}\mathbf{w}}]$ as inputs to the critic network. In domains without vision sensors, we simply use $\mathbf{o}_t^{\mathbf{1}\mathbf{o}\mathbf{w}}$ as \mathbf{h}_t without encoding the low-dimensional observations.

Coarse-to-fine Q-Network Coarse-to-fine Q-Network (CQN; Seo et al. 2024) is a value-based RL algorithm for continuous control that trains RL agents to zoom-into the continuous action space in a coarse-to-fine manner. In particular, CQN iterates the procedures of (i) discretizing the continuous action space into multiple bins and (ii) selecting the bin with the highest Q-value to further discretize. This reformulates the continuous control problem as a multi-level discrete control problem, allowing for the use of ideas from sample-efficient value-based RL algorithms (Mnih et al., 2015; Silver et al., 2017; Schrittwieser et al., 2020), designed to be used with discrete actions, for continuous control.

Formally, let \mathbf{a}_t^l be an action at level l with \mathbf{a}_t^0 being the zero vector¹. We then define the coarse-to-fine critic to consist of multiple Q-networks which compute Q-values for actions at each level \mathbf{a}_t^l , given the features \mathbf{h}_t and actions from the previous level \mathbf{a}_t^{l-1} , as follows:

$$Q_{\theta}^{l}(\mathbf{h}_{t}, \mathbf{a}_{t}^{l}, \mathbf{a}_{t}^{l-1}) \text{ for } l \in \{1, \dots, L\}$$

$$\tag{1}$$

131 We optimize each Q-network at level l with the following objective:

$$\mathcal{L}^{l} = \left(Q_{\theta}^{l}(\mathbf{h}_{t}, \mathbf{a}_{t}^{l}, \mathbf{a}_{t}^{l-1}) - r_{t+1} - \gamma \max_{a'} Q_{\overline{\theta}}^{l}(\mathbf{h}_{t+1}, a', \pi^{l}(\mathbf{h}_{t+1})) \right),$$
(2)

where $\bar{\theta}$ are delayed parameters for a target network (Polyak & Juditsky, 1992) and π^l is a policy that outputs the action \mathbf{a}_t^l at each level l via the inference steps with our critic, *i.e.*, $\pi^l(\mathbf{h}_t) = \mathbf{a}_t^l$. Specifically, to output actions at time step t with the critic, CQN first initializes constants $a_t^{1 \circ w}$ and a_t^{high} with -1 and 1. Then the following steps are repeated for $l \in \{1, ..., L\}$:

- Step 1 (Discretization): Discretize an interval $[a_t^{low}, a_t^{high}]$ into B uniform intervals, and each of these intervals become an action space for Q_{θ}^l
- Step 2 (Bin selection): Find a bin with the highest Q-value and set a_t^l to the centroid of the bin.
- Step 3 (Zoom-in): Set a_t^{low} and a_t^{high} to the minimum and maximum of the selected bin, which intuitively can be seen as zooming-into each bin.

We then use the last level's action \mathbf{a}_t^L as the action at time step t. For more details, including the inference procedure for computing Q-values, we refer readers to Appendix B.

3 COARSE-TO-FINE Q-NETWORK WITH ACTION SEQUENCE

We present Coarse-to-fine Q-Network with Action Sequence (CQN-AS), a value-based RL algorithm that learns a critic network that outputs Q-values for *a sequence of actions* $\mathbf{a}_{t:t+K} = {\mathbf{a}_t, ..., \mathbf{a}_{t+K-1}}$ for a given observation \mathbf{o}_t . Our main motivation for this design comes from one of the key ideas in recent behavior-cloning (BC) approaches, *i.e.*, predicting *action sequences*, which helps resolve ambiguity when approximating noisy, multi-modal distributions of expert demonstrations (Zhao et al., 2023; Chi et al., 2023). Similarly, by explicitly learning Q-values of both current and future actions from the given state, our approach aims to mitigate the challenge of learning Q-values with noisy trajectories from exploratory behaviors or human-collected demonstrations.

This section describes how we design our critic network with action sequence (see Section 3.1) and how we utilize action sequence outputs to control robots at each time step (see Section 3.2). The overview of our algorithm is available in Figure 2.

¹For simplicity, we describe CQN and CQN-**AS** with a single-dimensional action in the main section. See Appendix B for full description with *N*-dimensional actions, which is straightforward but requires more indices.

162 3.1 COARSE-TO-FINE CRITIC WITH ACTION SEQUENCE

167

171

172

173 174 175

192

193

205

Our key idea is to design a critic network to explicitly learn Q-values for current action and future actions from the current time step t, *i.e.*, $\{Q(\mathbf{o}_t, \mathbf{a}_t), Q(\mathbf{o}_t, \mathbf{a}_{t+1}), ..., Q(\mathbf{o}_t, \mathbf{a}_{t+K-1})\}$, to enable the critic to understand the consequence of executing a series of actions from the given state.

Formulation and objective Let $\mathbf{a}_{t:t+K}^{l} = {\mathbf{a}_{t}^{l}, ..., \mathbf{a}_{t+K-1}^{l}}$ be an action sequence at level l and $\mathbf{a}_{t:t+K}^{0}$ be a zero vector. We design our coarse-to-fine critic network to consist of multiple Q-networks that compute Q-values for each action at sequence step k and level l:

$$Q_{\theta}^{l,k}(\mathbf{h}_t, \mathbf{a}_{t+k-1}^l, \mathbf{a}_{t:t+K}^{l-1}) \text{ for } l \in \{1, ..., L\} \text{ and } k \in \{1, ..., K\}$$
(3)

We optimize our critic network with the following objective:

$$\sum_{k} \sum_{l} \left(Q_{\theta}^{l,k}(\mathbf{h}_{t}, \mathbf{a}_{t+k-1}^{l}, \mathbf{a}_{t:t+K}^{l-1}) - r_{t+1} - \gamma \max_{a'} Q_{\bar{\theta}}^{l,k}(\mathbf{h}_{t+1}, a', \pi_{K}^{l}(\mathbf{h}_{t+1})) \right)^{2}, \quad (4)$$

where π_K^l is an action sequence policy that outputs the action sequence $\mathbf{a}_{t:t+K}^l$. In practice, we compute Q-values for all sequence step $k \in \{1, ..., K\}$ in parallel, which is possible because Q-values for future actions depend only on current features \mathbf{h}_t but not on Q-values for previous actions. We find this simple design, with independence across action sequence, works well even on challenging humanoid control tasks with high-dimensional action spaces (Sferrazza et al., 2024). We expect our idea can be strengthened by exploiting the sequential structure, *i.e.*, Q-values at subsequent steps depend on previous Q-values (Metz et al., 2017; Chebotar et al., 2023), but we leave it as future work.

183 **Architecture** We implement our critic network to initially extract features for each sequence step k184 and aggregate features from multiple steps with a recurrent network (see Figure 2). This architecture 185 is often helpful in cases where a single-step action is already high-dimensional so that concatenating 186 them make inputs too high-dimensional. Specifically, let \mathbf{e}_k denote an one-hot encoding for k. At each level l, we construct features for each sequence step k as $\mathbf{h}_{t,k}^{l} = [\mathbf{h}_{t}, \mathbf{a}_{t+k-1}^{l-1}, \mathbf{e}_{k}]$. We then 187 188 encode each $\mathbf{h}_{t,k}^{l}$ with a shared MLP network and process them through GRU (Cho et al., 2014) to obtain $\mathbf{s}_{t,k}^{l} = f_{\theta}^{\text{GRU}}(f_{\theta}^{\text{MLP}}(\mathbf{h}_{t,1}^{l}), ..., f_{\theta}^{\text{MLP}}(\mathbf{h}_{t,k}^{l}))$. We then use a shared projection layer to map each $\mathbf{s}_{t,k}^{l}$ into Q-values at each sequence step k, *i.e.*, $Q_{\theta}^{l,k}(\mathbf{o}_{t}, \mathbf{a}_{t+k-1}^{l}, \mathbf{a}_{t:t+K}^{l-1}) = f_{\theta}^{\text{proj}}(\mathbf{s}_{t,k}^{l})$. 189 190 191

3.2 ACTION EXECUTION AND TRAINING DETAILS

While the idea of using action sequence is simple, there are two important yet small details for effectively training RL agents with action sequence: (i) how we execute actions at each time step to control robots and (ii) how we store training data and sample batches for training.

Executing action with temporal ensemble With the policy that outputs an action sequence $\mathbf{a}_{t:t+K}$, one important question is how to execute actions at time step $i \in \{t, ..., t + K - 1\}$. For this, we use the idea of Zhao et al. (2023) that utilizes *temporal ensemble*, which computes $\mathbf{a}_{t:t+K}$ every time step, saves it to a buffer, and executes a weighted average $\sum_i w_i \mathbf{a}_{t-i} / \sum w_i$ where $w_i = \exp(-m * i)$ denotes a weight that assigns higher value to more recent actions. We find this scheme outperforms the alternative of computing $\mathbf{a}_{t:t+K}$ every K steps and executing each action for subsequent K steps on most tasks we considered, except on several tasks that need reactive control.

Storing training data from environment interaction When storing samples from online environment interaction, we store a transition $(o_t, \hat{a}_t, r_{t+1}, o_{t+1})$ where \hat{a}_t denotes an action executed at time step t. For instance, if we use temporal ensemble for action execution, \hat{a}_t is a weighted average of action outputs obtained from previous K time steps.

Sampling training data from a replay buffer When sampling training data from the replay buffer, we sample a transition with action sequence, *i.e.*, $(o_t, \hat{a}_{t:t+K}, r_{t+1}, o_{t+1})$. If we sample time step t near the end of episode so that we do not have enough data to construct a full action sequence, we fill the action sequence with *null* actions. In particular, in position control where we specify the position of joints or end effectors, we repeat the action from the last step so that the agent learns not to change the position. On the other hand, in torque control where we specify the force to apply to joints, we set the action after the last step to zero so that agent learns to not to apply force.

225

226

227 228

229 230

231

232

233

234

235

236 237

238

239 240

241

242



Figure 3: Examples of robotic tasks. We study CON-AS on 53 robotic tasks spanning mobile bimanual manipulation, whole-body control, and tabletop manipulation tasks from BiGym (Chernyadev et al., 2024), HumanoidBench (Sferrazza et al., 2024), and RLBench (James et al., 2020).

4 EXPERIMENT

We study CON-AS on 53 robotic tasks spanning mobile bi-manual manipulation, whole-body control, and tabletop manipulation tasks from BiGym (Chernyadev et al., 2024), HumanoidBench (Sferrazza et al., 2024), and RLBench (James et al., 2020) environments (see Figure 3 for examples of robotic tasks). These tasks with sparse and dense rewards, with or without vision sensors, and with or without demonstrations, allow for evaluating the capabilities and limitations of our algorithm. In particular, our experiments are designed to investigate the following questions:

- Can CQN-AS quickly match the performance of a recent BC algorithm (Zhao et al., 2023) and surpass it through online learning? How does CQN-AS compare to previous model-free RL algorithms (Haarnoja et al., 2018; Yarats et al., 2022; Seo et al., 2024)?
- What is the contribution of each component in CQN-AS?
- Under which conditions is CQN-AS effective? When does CQN-AS fail?

243 Baselines for fine-grained control tasks with demonstrations For tasks that need high-precision 244 control, e.g., manipulation tasks from BiGym and RLBench, we consider model-free RL baselines 245 that aim to learn deterministic policies, as we find that stochastic policies struggle to solve such 246 fine-grained control tasks. Specifically, we consider (i) Coarse-to-fine Q-Network (CQN; Seo et al. 2024), a value-based RL algorithm that learns to zoom-into continuous action space in a coarse-to-fine 247 manner, and (ii) DrQ-v2+, an optimized demo-driven variant of a model-free actor-critic algorithm 248 DrQ-v2 (Yarats et al., 2022) that uses a deterministic policy algorithm and data augmentation. We 249 further consider (iii) Action Chunking Transformer (ACT; Zhao et al. 2023), a BC algorithm that 250 trains a transformer (Vaswani et al., 2017) policy to predict action sequence and utilizes temporal 251 ensemble for executing actions, as our highly-optimized BC baseline. 252

253 **Baselines for whole-body control tasks with dense reward** For locomotion tasks with dense 254 reward, we consider (i) Soft Actor-Critic (SAC; Haarnoja et al. 2018), a model-free actor-critic RL 255 algorithm that maximizes action entropy, and (ii) Coarse-to-fine Q-Network (CQN; Seo et al. 2024). 256 Moreover, although it is **not** the goal of this paper to compare against model-based RL algorithms, 257 we also consider two model-based baselines: (iii) DreamerV3 (Hafner et al., 2023), a model-based 258 RL algorithm that learns a latent dynamics model and a policy using imagined trajectories and (iv) TD-MPC2 (Hansen et al., 2024), a model-based RL algorithm that learns a latent dynamics model 259 and utilizes local trajectory optimization in imagined latent trajectories. 260

261 **Implementation details** For training with expert demonstrations, we follow the setup of Seo et al. 262 (2024). Specifically, we keep a separate replay buffer that stores demonstrations and sample half 263 of training data from demonstrations. We also relabel successful online episodes as demonstrations 264 and store them in the demonstration replay buffer. For CQN-AS, we use an auxiliary BC loss from 265 Seo et al. (2024) based on large margin loss (Hester et al., 2018). For actor-critic baselines, we 266 use an auxiliary BC loss that minimizes L2 loss between the policy outputs and expert actions. We 267 implement CQN-AS based on a publicly available source code of CQN² based on PyTorch (Paszke 268 et al., 2019). We will release the full source code upon publication. 269

²https://github.com/younggyoseo/CQN



Figure 4: **BiGym results** on 25 sparsely-rewarded mobile bi-manual manipulation tasks (Chernyadev et al., 2024). All experiments are initialized with 17 to 60 *human-collected* demonstrations, and RL methods are trained with an auxiliary BC objective. On many of the challenging long-horizon tasks, CQN-**AS** quickly learns to match the performance of ACT (Zhao et al., 2023) and surpass it through online learning, while other RL baselines fail to effectively accelerate training with noisy human-collected demonstrations. We report the success rate over 25 episodes. The solid line and shaded regions represent the mean and confidence intervals, respectively, across 4 runs.

299

300

301

302

303

304

306

307

308

309

4.1 **BIGYM EXPERIMENTS**

We study CQN-AS on mobile bi-manual manipulation tasks from BiGym (Chernyadev et al., 2024). BiGym's *human-collected* demonstrations are often noisy and multi-modal, posing challenges to RL algorithms which should effectively leverage demonstrations for solving sparsely-rewarded tasks.

310 311

Setup Because we find that not all demonstrations from BiGym benchmark can be successfully 312 replayed³, we replay all the demonstrations and only use the successful ones as demonstrations. 313 We do not discard ones that fail to be replayed, but we use them as training data with zero reward 314 because they can still be useful as failure experiences. To avoid training with too few demon-315 strations, we exclude the tasks where the ratio of successful demonstrations is below 50%. This 316 leaves us with 25 tasks, each with 17 to 60 demonstrations. For visual observations, we use RGB 317 observations with 84×84 resolution from head, left_wrist, and right_wrist cameras. We also 318 use low-dimensional proprioceptive states from proprioception, proprioception_grippers, and 319 proprioception_floating_base sensors. We use (i) absolute joint position control action mode and (ii) floating base that replaces locomotion with classic controllers. We use the same set of 320 hyperparameters for all the tasks, in particular, we use action sequence of length 16. More details on 321 BiGym experiments are available in Appendix A. 322

323

³We use demonstrations available at the date of Oct 1st with the commit 018f8b2.

339

340

341

342

343

344

345



Figure 5: HumanoidBench results on 8 densely-rewarded humanoid control tasks (Sferrazza et al., 2024). All the experiments start from scratch and RL methods do not have an auxiliary BC objective. CQN-AS significantly outperforms other model-free RL baselines on most tasks. CQN-AS often achieves competitive performance to model-based RL baselines, which is intriguing but not the main goal of this paper. For CQN-AS and CQN, we report the results aggregated over 4 runs. For other baselines, we report the results aggregated over 3 runs available from public website. The solid line and shaded regions represent the mean and confidence intervals.

346 **Comparison to baselines** Figure 4 shows the experimental results on BiGym benchmark. We find 347 that CQN-AS quickly matches the performance of ACT and outperforms it through online learning 348 on 20/25 tasks, while other RL algorithms fail to do so especially on challenging long-horizon tasks 349 such as Move Plate and Saucepan To Hob. A notable result here is that CQN-AS enables solving 350 these challenging BiGym tasks while other RL baselines, *i.e.*, CQN and DrQ-v2+, completely fail 351 as they achieve 0% success rate. This result highlights the capability of CQN-AS to accelerate RL 352 training from noisy, multi-modal demonstrations collected by humans. 353

Limitation However, we find that CQN-AS struggles to achieve meaningful success rate on some of the long-horizon tasks that require interaction with delicate objects such as cup or cutlery, leaving room for future work to incorporate advanced vision encoders (He et al., 2016; 2022) or critic architectures (Kapturowski et al., 2023; Chebotar et al., 2023; Springenberg et al., 2024).

4.2 HUMANOIDBENCH EXPERIMENTS

To show that CON-AS can be generally applicable to tasks without demonstrations, we study CON-AS 361 on densely-rewarded humanoid control tasks from HumanoidBench (Sferrazza et al., 2024). 362

354

355

356

357

358 359

360

Setup For HumanoidBench, we follow a standard setup that trains RL agents from scratch, which is 364 also used in original benchmark (Sferrazza et al., 2024). Specifically, we use low-dimensional states 365 consisting of proprioception and privileged task information as inputs. For tasks, we simply select the 366 first 8 locomotion tasks in the benchmark. Following the original benchmark that trains RL agents 367 for environment steps that roughly requires 48 hours of training, we report the results of CQN-AS 368 and CON for 7 million steps. For baselines, we use the results available from the public repository, 369 which are evaluated on tasks with dexterous hands, and we also evaluate our algorithm on tasks with 370 hands. We use the same set of hyperparameters for all the tasks, in particular, we use action sequence 371 of length 4. More details on HumanoidBench experiments are available in Appendix A.

372

373 **Comparison to model-free RL baselines** Figure 5 shows the results on on HumanoidBench. We 374 find that, by learning the critic network with action sequence, CQN-AS outperforms other model-free 375 RL baselines, *i.e.*, CQN and SAC, on most tasks. In particular, the difference between CQN-AS and baselines becomes larger as the task gets more difficult, e.g., baselines fail to achieve high episode 376 return on Walk and Run tasks but CQN-AS achieves strong performance. This result shows that our 377 idea of using action sequence can be applicable to generic setup without demonstrations.



Figure 6: **RLBench results** on 20 sparsely-rewarded tabletop manipulation tasks from RLBench (James et al., 2020). All experiments are initialized with 100 *synthetic* demonstrations generated via motion-planning and RL methods are trained with an auxiliary BC objective. *As expected*, with synthetic demonstrations, CQN-AS achieves *similar* performance to CQN on most tasks. However, CQN-AS often significantly outperforms baselines on several challenging, long-horizon tasks such as Open Oven. We report the success rate over 25 episodes. The solid line and shaded regions represent the mean and confidence intervals, respectively, across 4 runs.

408
 409
 409
 409
 409
 409
 410
 410
 410
 410
 410
 411
 411
 412
 412
 413
 413
 414
 415
 416
 417
 418
 418
 419
 419
 410
 410
 410
 411
 411
 411
 412
 412
 413
 413
 414
 415
 415
 416
 417
 418
 418
 419
 419
 410
 410
 410
 411
 411
 412
 412
 413
 414
 415
 416
 417
 418
 418
 419
 419
 410
 410
 411
 411
 412
 412
 413
 414
 415
 415
 416
 417
 418
 418
 419
 419
 410
 410
 411
 411
 412
 412
 413
 414
 415
 415
 416
 417
 418
 418
 419
 419
 419
 410
 410
 411
 411
 412
 412
 413
 414
 415
 415
 416
 417
 418
 418
 419
 419
 419
 419
 410
 410
 411
 411
 412
 411
 412
 412
 413

414 415

416

417

4.3 RLBENCH EXPERIMENTS

To investigate whether CQN-AS can also be effective in leveraging *clean* demonstrations, we study CQN-AS on RLBench (James et al., 2020) with synthetic demonstrations.

418 419

 Setup For RLBench experiments, we use the official CQN implementation for collecting demonstrations and reproducing the baseline results on the same set of tasks. Specifically, we use RGB observations with 84×84 resolution from front, wrist, left_shoulder, and right_shoulder cameras. We also use low-dimensional proprioceptive states consisting of 7-dimensional joint positions and a binary value for gripper open. We use 100 demonstrations and delta joint position control action mode. We use the same set of hyperparameters for all the tasks, in particular, we use action sequence of length 4. More details on RLBench experiments are available in Appendix A.

427 CQN-AS is also effective with *clean* demonstrations Because RLBench provides synthetic *clean* 428 demonstrations, as we expected, Figure 6 shows that CQN-AS achieves *similar* performance to CQN
 429 on many of the tasks, except 2/25 tasks where it hurts the performance. But we still find that CQN-AS
 430 achieves quite superior performance to CQN on some challenging long-horizon tasks such as Open
 431 Oven or Take Plate Off Colored Dish Rack. These results, along with results from BiGym and
 HumanoidBench, show that CQN-AS can be used in various benchmark with different characteristics.



Figure 7: **Ablation studies and analysis** on the effect of (a) action sequence, (b) RL objective, and (c) temporal ensemble. (d) We also provide results on locomotion tasks from DeepMind Control Suite (Tassa et al., 2020), where CQN-**AS** fails to improve performance. The solid line and shaded regions represent the mean and confidence intervals, respectively, across 4 runs.

4.4 Ablation Studies, Analysis, Failure Cases

Effect of action sequence length Figure 7a shows the performance of CQN-AS with different action sequence lengths on two BiGym tasks. We find that training the critic network with longer action sequences improves and stabilizes performance.

RL objective is crucial for strong performance Figure 7b shows the performance of CQN-AS without RL objective that trains the model only with BC objective on successful demonstrations. We find this baseline significantly underperforms CQN-AS, which shows that RL objective is indeed enabling the agent to learn from online trial-and-error experiences.

Effect of temporal ensemble Figure 7c shows that performance largely degrades without temporal ensemble on Saucepan To Hop as temporal ensemble induces a smooth motion and thus improves performance in fine-grained control tasks. But we also find that temporal ensemble can be harmful on Reach Target Single. We hypothesize this is because temporal ensemble often makes it difficult to refine behaviors based on recent visual observations. Nonetheless, we use temporal ensemble for all the tasks as we find it helps on most tasks and we aim to use the same set of hyperparameters.

Failure case: Torque control Figure 7d shows that CQN-AS underperforms CQN on locomotion tasks with torque control. We hypothesize this is because a sequence of joint positions usually has a semantic meaning in joint spaces, making it easier to learn with, when compared to learning how to apply a sequence of torques. Addressing this failure case is an interesting future direction.

5 RELATED WORK

Behavior cloning with action sequence Recent behavior cloning approaches have shown that predicting a sequence of actions enables the policy to effectively imitate noisy expert trajectories and helps in dealing with idle actions from human pauses during data collection (Zhao et al., 2023; Chi et al., 2023). In particular, Zhao et al. (2023) train a transformer model (Vaswani et al., 2017) that predicts action sequence and Chi et al. (2023) train a denoising diffusion model (Ho et al., 2020) that approximates the action distributions. This idea has been extended to multi-task setup (Bharadhwaj et al., 2024), mobile manipulation (Fu et al., 2024b) and humanoid control (Fu et al., 2024a). Our work is inspired by this line of work and proposed to learn RL agents with action sequence.

486 **Reinforcement learning with action sequence** In the context of reinforcement learning, Medini & 487 Shrivastava (2019) proposed to pre-compute frequent action sequences from expert demonstrations 488 and augment the action space with these sequences. However, this idea introduces additional 489 complexity and is not scalable to setups without demonstrations. One recent work relevant to ours 490 is Saanum et al. (2024) that encourages a sequence of actions from RL agents to be predictable and smooth. But this differs from our work in that it uses the concept of action sequence only for 491 computing the penalty term. Recently, Ankile et al. (2024) point out that RL with action sequence 492 is challenging and instead proposes to use RL for learning a single-step policy that corrects action 493 sequence predictions from BC. In contrast, our work shows that training RL agents with action 494 sequence is feasible and leads to improved performance compared to prior RL algorithms. 495

496 497

498

6 CONCLUSION

We presented Coarse-to-fine Q-Network with Action Sequence (CQN-AS), a value-based RL algorithm that trains a critic network that outputs Q-values over action sequences. Extensive experiments in benchmarks with various setups show that our idea not only improves the performance of the base algorithm but also allows for solving complex tasks where prior RL algorithms completely fail.

We believe our work will be strong evidence that shows RL can realize its promise to develop robots
that can continually improve through online trial-and-error experiences, surpassing the performance
of BC approaches. We are excited about future directions, including real-world RL with humanoid
robots, incorporating advanced critic architectures (Kapturowski et al., 2023; Chebotar et al., 2023;
Springenberg et al., 2024), bootstrapping RL agents from imitation learning (Hu et al., 2023; Xing
et al., 2024) or offline RL (Nair et al., 2020; Lee et al., 2021), extending the idea to recent model-based
RL approaches (Hafner et al., 2023; Hansen et al., 2024), to name but a few.

510

511 REPRODUCIBILITY STATEMENT 512

We have provided details required to implement our algorithm and reproduce the results in Section 4 and Appendix A. We will release the full source code upon publication.

515 516

523

524

525 526

527

528

529

530

531

532

533

517 REFERENCES

- Lars Ankile, Anthony Simeonov, Idan Shenfeld, Marcel Torne, and Pulkit Agrawal. From imitation to refinement–residual rl for precise visual assembly. *arXiv preprint arXiv:2407.16677*, 2024.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
 - Marc G Bellemare, Will Dabney, and Rémi Munos. A distributional perspective on reinforcement learning. In *International Conference on Machine Learning*, 2017.
 - Homanga Bharadhwaj, Jay Vakil, Mohit Sharma, Abhinav Gupta, Shubham Tulsiani, and Vikash Kumar. Roboagent: Generalization and efficiency in robot manipulation via semantic augmentations and action chunking. In 2024 IEEE International Conference on Robotics and Automation (ICRA), 2024.
 - Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning*, 2023a.
- Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. In *Robotics: Science and Systems*, 2023b.
- Yevgen Chebotar, Quan Vuong, Karol Hausman, Fei Xia, Yao Lu, Alex Irpan, Aviral Kumar, Tianhe
 Yu, Alexander Herzog, Karl Pertsch, et al. Q-transformer: Scalable offline reinforcement learning via autoregressive q-functions. In *Conference on Robot Learning*, 2023.

540 541 542	Nikita Chernyadev, Nicholas Backshall, Xiao Ma, Yunfan Lu, Younggyo Seo, and Stephen James. Bigym: A demo-driven mobile bi-manual manipulation benchmark. In <i>Conference on Robot Learning</i> , 2024.
544 545 546	Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In <i>Robotics: Science and Systems</i> , 2023.
547 548 549 550	Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. <i>arXiv preprint arXiv:1406.1078</i> , 2014.
551 552	Zipeng Fu, Qingqing Zhao, Qi Wu, Gordon Wetzstein, and Chelsea Finn. Humanplus: Humanoid shadowing and imitation from humans. In <i>Conference on Robot Learning</i> , 2024a.
553 554 555	Zipeng Fu, Tony Z Zhao, and Chelsea Finn. Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. In <i>Conference on Robot Learning</i> , 2024b.
556 557 558	Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. <i>arXiv preprint arXiv:1812.05905</i> , 2018.
559 560 561	Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. <i>arXiv preprint arXiv:2301.04104</i> , 2023.
562 563	Nicklas Hansen, Hao Su, and Xiaolong Wang. Td-mpc2: Scalable, robust world models for continuous control. In <i>International Conference on Learning Representations</i> , 2024.
565 566 567	Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In <i>Proceedings of the IEEE conference on computer vision and pattern recognition</i> , 2016.
568 569 570	Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In <i>Proceedings of the IEEE/CVF conference on computer vision and pattern recognition</i> , 2022.
572 573	Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). arXiv preprint arXiv:1606.08415, 2016.
574 575 576 577	Alexander Herzog, Kanishka Rao, Karol Hausman, Yao Lu, Paul Wohlhart, Mengyuan Yan, Jessica Lin, Montserrat Gonzalez Arenas, Ted Xiao, Daniel Kappler, et al. Deep rl at scale: Sorting waste in office buildings with a fleet of mobile manipulators. <i>arXiv preprint arXiv:2305.03270</i> , 2023.
578 579 580	Todd Hester, Matej Vecerik, Olivier Pietquin, Marc Lanctot, Tom Schaul, Bilal Piot, Dan Horgan, John Quan, Andrew Sendonaris, Ian Osband, et al. Deep q-learning from demonstrations. In <i>Proceedings of the AAAI conference on artificial intelligence</i> , 2018.
581 582 583	Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. Advances in neural information processing systems, 2020.
584 585	Hengyuan Hu, Suvir Mirchandani, and Dorsa Sadigh. Imitation bootstrapped reinforcement learning. arXiv preprint arXiv:2311.02198, 2023.
586 587 588	Stephen James, Marc Freese, and Andrew J Davison. Pyrep: Bringing v-rep to deep robot learning. arXiv preprint arXiv:1906.11176, 2019.
589 590 591	Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J Davison. Rlbench: The robot learning benchmark & learning environment. <i>IEEE Robotics and Automation Letters</i> , 5(2):3019–3026, 2020.
592 593	Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. <i>Artificial intelligence</i> , 1998.

594 595 596	Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, et al. Scalable deep reinforcement learning for vision-based robotic manipulation. In <i>Conference on robot learning</i> , 2018.
597 598 599 600	Steven Kapturowski, Víctor Campos, Ray Jiang, Nemanja Rakićević, Hado van Hasselt, Charles Blundell, and Adrià Puigdomènech Badia. Human-level atari 200x faster. In <i>International Conference on Learning Representations</i> , 2023.
601 602 603	Seunghyun Lee, Younggyo Seo, Kimin Lee, Pieter Abbeel, and Jinwoo Shin. Offline-to-online reinforcement learning via balanced replay and pessimistic q-ensemble. In <i>Conference on Robot Learning</i> , 2021.
605 606 607	Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In <i>International Conference on Learning Representations</i> , 2016.
608 609	Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In International Confer- ence on Learning Representations, 2019.
610 611 612 613	Tharun Medini and Anshumali Shrivastava. Mimicking actions is a good strategy for beginners: Fast reinforcement learning with expert action sequences, 2019. URL https://openreview.net/forum?id=HJfxbhR9KQ.
614 615	Luke Metz, Julian Ibarz, Navdeep Jaitly, and James Davidson. Discrete sequential prediction of continuous actions for deep rl. <i>arXiv preprint arXiv:1705.05035</i> , 2017.
617 618 619	Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. <i>Nature</i> , 2015.
620 621	Ashvin Nair, Abhishek Gupta, Murtaza Dalal, and Sergey Levine. Awac: Accelerating online reinforcement learning with offline datasets. <i>arXiv preprint arXiv:2006.09359</i> , 2020.
622 623 624 625 626	Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. <i>Advances in neural information processing systems</i> , 32, 2019.
627 628	Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. <i>SIAM journal on control and optimization</i> , 1992.
629 630 631	Dean A Pomerleau. Alvinn: An autonomous land vehicle in a neural network. In Advances in neural information processing systems, 1988.
632 633 634	Eric Rohmer, Surya PN Singh, and Marc Freese. V-rep: A versatile and scalable robot simulation framework. In <i>IEEE/RSJ international conference on intelligent robots and systems</i> , 2013.
635 636	Tankred Saanum, Noémi Éltető, Peter Dayan, Marcel Binz, and Eric Schulz. Reinforcement learning with simple sequence priors. <i>Advances in Neural Information Processing Systems</i> , 2024.
637 638 639 640	Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. <i>Nature</i> , 2020.
641 642	Frank Sehnke, Christian Osendorfer, Thomas Rückstieß, Alex Graves, Jan Peters, and Jürgen Schmidhuber. Parameter-exploring policy gradients. <i>Neural Networks</i> , 23(4):551–559, 2010.
643 644 645	Younggyo Seo, Jafar Uruç, and Stephen James. Continuous control with coarse-to-fine reinforcement learning. In <i>Conference on Robot Learning</i> , 2024.
646 647	Tim Seyde, Peter Werner, Wilko Schwarting, Igor Gilitschenski, Martin Riedmiller, Daniela Rus, and Markus Wulfmeier. Solving continuous control via q-learning. In <i>International Conference on Learning Representations</i> , 2023.

648 649 650	Carmelo Sferrazza, Dun-Ming Huang, Xingyu Lin, Youngwoon Lee, and Pieter Abbeel. Humanoid- bench: Simulated humanoid benchmark for whole-body locomotion and manipulation. In <i>Robotics:</i> <i>Science and Systems</i> , 2024.
651 652 653 654	David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. <i>nature</i> , 2017.
655 656 657 658	Jost Tobias Springenberg, Abbas Abdolmaleki, Jingwei Zhang, Oliver Groth, Michael Bloesch, Thomas Lampe, Philemon Brakel, Sarah Bechtle, Steven Kapturowski, Roland Hafner, et al. Offline actor-critic reinforcement learning scales to large models. In <i>International Conference on</i> <i>Machine Learning</i> , 2024.
659 660	Richard S Sutton and Andrew G Barto. Reinforcement learning: An introduction. MIT press, 2018.
661 662 663	Yuval Tassa, Saran Tunyasuvunakool, Alistair Muldal, Yotam Doron, Siqi Liu, Steven Bohez, Josh Merel, Tom Erez, Timothy Lillicrap, and Nicolas Heess. dm_control: Software and tasks for continuous control. <i>arXiv preprint arXiv:2006.12983</i> , 2020.
664 665 666	Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, et al. Octo: An open-source generalist robot policy. In <i>Robotics: Science and Systems</i> , 2024.
668 669	Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE, 2012.
670 671 672	Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In <i>Advances in Neural Information Processing Systems</i> , 2017.
674 675 676	Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Hasselt, Marc Lanctot, and Nando Freitas. Dueling network architectures for deep reinforcement learning. In <i>International conference on machine learning</i> , 2016.
677 678 679	Jiaxu Xing, Angel Romero, Leonard Bauersfeld, and Davide Scaramuzza. Bootstrapping reinforce- ment learning with imitation for vision-based agile flight. In <i>Conference on Robot Learning</i> , 2024.
680 681 682 683	Denis Yarats, Rob Fergus, Alessandro Lazaric, and Lerrel Pinto. Mastering visual continuous control: Improved data-augmented reinforcement learning. In <i>International Conference on Learning</i> <i>Representations</i> , 2022.
684 685 686 687 688 689 690 691 692 693 694 695 696 697 698 699	Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. In <i>Robotics: Science and Systems</i> , 2023.
700	

702 A EXPERIMENTAL DETAILS

BiGym BiGym⁴ (Chernyadev et al., 2024) is built upon MuJoCo (Todorov et al., 2012). We use
Unitree H1 with two parallel grippers. We find that demonstrations available in the recent version of
BiGym are not all successful. Therefore we adopt the strategy of replaying all the demonstrations and
only use the successful ones as demonstrations. instead of discarding the failed demonstrations, we
still store them in a replay buffer as failure experiences. To avoid training with too few demonstrations,
we exclude the tasks where the ratio of successful demonstrations is below 50%. Table 1 shows the
list of 25 sparsely-rewarded mobile bi-manual manipulation tasks used in our experiments.

Table 1: BiGym tasks with their maximum episode length and number of successful demonstrations.

Task	Length	Demos	Task	Length	Demos
Move Plate	300	51	Cupboards Close All	620	53
Move Two Plates	550	30	Reach Target Single	100	30
Saucepan To Hob	440	28	Reach Target Multi Modal	100	60
Sandwich Flip	620	34	Reach Target Dual	100	50
Sandwich Remove	540	24	Dishwasher Close	375	44
Dishwasher Load Plates	560	17	Wall Cupboard Open	300	44
Dishwasher Load Cups	750	58	Drawers Open All	480	45
Dishwasher Unload Cutlery	620	29	Wall Cupboard Close	300	60
Take Cups	420	32	Dishwasher Open Trays	380	57
Put Cups	425	43	Drawers Close All	200	59
Flip Cup	550	45	Drawer Top Open	200	40
Flip Cutlery	500	43	Drawer Top Close	120	51
Dishwasher Close Trays	320	62	-		

HumanoidBench HumanoidBench⁵ (Sferrazza et al., 2024) is built upon MuJoCo (Todorov et al., 2012). We use Unitree H1 with two dexterous hands. We consider the first 8 locomotion tasks in the benchmark: Stand, Walk, Run, Reach, Hurdle, Crawl, Maze, Sit Simple. We use proprioceptive states and privileged task information instead of visual observations. Unlike BiGym and RLBench experiments, we do not utilize dueling network (Wang et al., 2016) and distributional critic (Bellemare et al., 2017) in HumanoidBench for faster experimentation.

RLBench RLBench⁶ (James et al., 2020) is built upon CoppeliaSim (Rohmer et al., 2013) and PyRep (James et al., 2019). We use a 7-DoF Franka Panda robot arm and a parallel gripper. Following the setup of Seo et al. (2024), we increase the velocity and acceleration of the arm by 2 times. For all experiments, we use 100 demonstrations generated via motion-planning. Table 2 shows the list of 20 sparsely-rewarded visual manipulation tasks used in our experiments.

Table 2: RLBench tasks with their maximum episode length used in our experiments.

Task	Length	Task	Length
Take Lid Off Saucepan	100	Put Books On Bookshelf	175
Open Drawer	100	Sweep To Dustpan	100
Stack Wine	150	Pick Up Cup	100
Toilet Seat Up	150	Open Door	125
Open Microwave	125	Meat On Grill	150
Open Oven	225	Basketball In Hoop	125
Take Plate Off Colored Dish Rack	150	Lamp On	100
Turn Tap	125	Press Switch	100
Put Money In Safe	150	Put Rubbish In Bin	150
Phone on Base	175	Insert Usb In Computer	100

754 ⁴https://github.com/chernyadev/bigym

⁵https://github.com/carlosferrazza/humanoid-bench

⁶https://github.com/stepjam/RLBench

761

790

791

Hyperparameters We use the same set of hyperparameters across the tasks in each domain. For
hyperparameters shared across CQN and CQN-AS, we use the same hyperparameters for both
algorithms for a fair comparison. We provide detailed hyperparameters for BiGym and RLBench
experiments in Table 3 and hyperparameters for HumanoidBench experiments in Table 4

Table 3: Hyperparameters for demo-driven vision-based experiments in BiGym and RLBench

762	Hyperperemeter	Value
763	11yper par ameter	value
764	Image resolution	$84 \times 84 \times 3$
765	Image augmentation	RandomShift (Yarats et al., 2022)
766	Frame stack	4 (BiGym) / 8 (RLBench)
767	CNN - Architecture	Conv (c=[32, 64, 128, 256], s=2, p=1)
768	MLP - Architecture	Linear (c=[512, 512, 64, 512, 512], bias=False) (BiGym) Linear (c=[64, 512, 512], bias=False) (RLBench)
769	CNN & MLP - Activation	SiLU (Hendrycks & Gimpel, 2016) and LayerNorm (Ba et al., 2016)
770	GRU - Architecture	GRU (c=[512], bidirectional=False)
771	Dueling network	True
772	C51 - Atoms	51
773	$C51 - v_{min}, v_{max}$	-2, 2
774	Action sequence	16 (BiGym) / 4 (RI Bench)
775	Temporal ensemble weight m	0.01
776	Levels	3
777	Bins	5
778	BC loss (\mathcal{L}_{BC}) scale	1.0
779	RL loss (\mathcal{L}_{RL}) scale	0.1
780	Relabeling as demonstrations	True
781	Data-driven action scaling	True
782	Action mode	Absolute Joint (BiGym), Delta Joint (RLBench)
783	Exploration noise Target critic undete ratio (π)	$\epsilon \sim \mathcal{N}(0, 0.01)$
79/	N-step return	1
704	Batch size	1 128 (BiGym) / 256 (RLBench)
785	Demo batch size	128 (BiGym) / 256 (RLBench)
786	Optimizer	AdamW (Loshchilov & Hutter, 2019)
787	Learning rate	5e-5
788	Weight decay	0.1
789		

Table 4: Hyperparameters for state-based experiments in HumanoidBench

Hyperparameter	Value		
MLP - Architecture	Linear (c=[512, 512], bias=False)		
CNN & MLP - Activation	SiLU (Hendrycks & Gimpel, 2016) and LayerNorm (Ba et al., 2016)		
GRU - Architecture	GRU (c=[512], bidirectional=False)		
Dueling network	False		
Action sequence	4		
Temporal ensemble weight m	0.01		
Levels	3		
Bins	5		
RL loss (\mathcal{L}_{RL}) scale	1.0		
Action mode	Absolute Joint		
Exploration noise	$\epsilon \sim \mathcal{N}(0, 0.01)$		
Target critic update ratio (τ)	1.0		
Target critic update interval (τ)	1000		
N-step return	3		
Batch size	128		
Optimizer	AdamW (Loshchilov & Hutter, 2019)		
Learning rate	5e-5		
Weight decay	0.1		

810 **Computing hardware** For all experiments, we use consumer-grade 11GB GPUs such as NVIDIA 811 GTX 1080Ti, NVIDIA Titan XP, and NVIDIA RTX 2080Ti with 11 or 12GB VRAM. With 2080Ti 812 GPU, each BiGym experiment with 100K environment steps take 9.5 hours, each RLBench experi-813 ment with 30K environment steps take 6.5 hours, and each HumanoidBench experiment with 7M 814 environment steps take 48 hours. We find that CQN-AS is around 33% slower than running CQN because larger architecture slows down both training and inference. 815

Baseline implementation For CQN (Seo et al., 2024) and DrQ-v2+ (Yarats et al., 2022), we use the implementation available from the official CQN implementation⁷. For ACT (Zhao et al., 818 2023), we use the implementation from RoboBase repository⁸. For SAC (Haarnoja et al., 2018), DreamerV3 (Hafner et al., 2023), and TD-MPC2 (Hansen et al., 2024), we use results provided in HumanoidBench⁹ repository (Sferrazza et al., 2024).

FULL DESCRIPTION OF CQN AND CQN-AS В

This section provides the formulation of CQN and CQN-AS with *n*-dimensional actions.

B.1 COARSE-TO-FINE Q-NETWORK

816

817

819

820

821 822

823 824

825 826

827 828

829

830 831

836

837 838

839 840

841

842 843

844

845

846 847

848 849

850

851 852

853 854 855

856

858

859

862

863

Let $a_t^{l,n}$ be an action at level l and dimension n and $\mathbf{a}_t^l = \{a_t^{l,1}, ..., a_t^{l,N}\}$ be actions at level l with \mathbf{a}_t^0 being zero vector. We then define coarse-to-fine critic to consist of multiple Q-networks:

$$Q_{\theta}^{l,n}(\mathbf{h}_t, a_t^{l,n}, \mathbf{a}_t^{l-1}) \text{ for } l \in \{1, ..., L\} \text{ and } n \in \{1, ..., N\}$$
(5)

We optimize the critic network with the following objective:

$$\sum_{n} \sum_{l} \left(Q_{\theta}^{l,n}(\mathbf{h}_{t}, a_{t}^{l,n}, \mathbf{a}_{t}^{l-1}) - r_{t+1} - \gamma \max_{a'} Q_{\bar{\theta}}^{l,n}(\mathbf{h}_{t+1}, a', \pi^{l}(\mathbf{h}_{t+1})) \right)^{2},$$
(6)

where $\bar{\theta}$ are delayed parameters for a target network (Polyak & Juditsky, 1992) and π^{l} is a policy that outputs the action \mathbf{a}_{t}^{l} at each level l via the inference steps with our critic, i.e., $\pi^{l}(\mathbf{h}_{t}) = \mathbf{a}_{t}^{l}$.

Action inference To output actions at time step t with the critic, CQN first initializes constants $a_t^{n,\text{low}}$ and $a_t^{n,\text{high}}$ with -1 and 1 for each n. Then the following steps are repeated for $l \in \{1, ..., L\}$:

- Step 1 (Discretization): Discretize an interval $[a_t^{n,low}, a_t^{n,high}]$ into B uniform intervals, and each of these intervals become an action space for $Q_{\theta}^{l,n}$
- Step 2 (Bin selection): Find the bin with the highest Q-value, set $a_t^{l,n}$ to the centroid of the selected bin, and aggregate actions from all dimensions to \mathbf{a}_t^l
- Step 3 (Zoom-in): Set $a_t^{n,low}$ and $a_t^{n,high}$ to the minimum and maximum of the selected bin, which intuitively can be seen as zooming-into each bin.

We then use the last level's action \mathbf{a}_t^L as the action at time step t.

Computing Q-values To compute Q-values for given actions \mathbf{a}_t , CQN first initializes constants $a_t^{n,\text{low}}$ and $a_t^{n,\text{high}}$ with -1 and 1 for each n. We then repeat the following steps for $l \in \{1, ..., L\}$:

- Step 1 (Discretization): Discretize an interval $[a_t^{n,low}, a_t^{n,high}]$ into B uniform intervals, and each of these intervals become an action space for $Q_{\theta}^{l,n}$
- Step 2 (Bin selection): Find the bin that contains input action \mathbf{a}_t , compute $a_t^{l,n}$ for the selected interval, and compute Q-values $Q_{\theta}^{l,n}(\mathbf{h}_t, a_t^{l,n}, \mathbf{a}_t^{l-1})$.
- Step 3 (Zoom-in): Set $a_t^{n, low}$ and $a_t^{n, high}$ to the minimum and maximum of the selected bin, which intuitively can be seen as zooming-into each bin.

```
We then use a set of Q-values \{Q_{\theta}^{l,n}(\mathbf{h}_t, \mathbf{a}_t^{l,n}, \mathbf{a}_t^{l-1})\}_{l=1}^L for given actions \mathbf{a}_t.
861
```

⁷https://github.com/younggyoseo/CQN

⁸https://github.com/robobase-org/robobase

⁹https://github.com/carlosferrazza/humanoid-bench

864 B.2 COARSE-TO-FINE CRITIC WITH ACTION SEQUENCE

Let $\mathbf{a}_{t:t+K}^{l} = {\mathbf{a}_{t}^{l}, ..., \mathbf{a}_{t+K-1}^{l}}$ be an action sequence at level l and $\mathbf{a}_{t:t+K}^{0}$ be zero vector. Our critic network consists of multiple Q-networks for each level l, dimension n, and sequence step k:

$$Q_{\theta}^{l,n,k}(\mathbf{h}_{t}, a_{t+k-1}^{l,n}, \mathbf{a}_{t:t+K}^{l-1}) \text{ for } l \in \{1, ..., L\}, \ n \in \{1, ..., N\} \text{ and } k \in \{1, ..., K\}$$
(7)

We optimize the critic network with the following objective:

$$\sum_{n} \sum_{l} \sum_{k} \left(Q_{\theta}^{l,n,k}(\mathbf{h}_{t}, a_{t}^{l,n}, \mathbf{a}_{t:t+K}^{l-1}) - r_{t+1} - \gamma \max_{a'} Q_{\bar{\theta}}^{l,n,k}(\mathbf{h}_{t+1}, a', \pi_{K}^{l}(\mathbf{h}_{t+1})) \right)^{2}, \quad (8)$$

where π_K^l is an action sequence policy that outputs the action sequence $\mathbf{a}_{t:t+K}^l$. In practice, we compute Q-values for all sequence step $k \in \{1, ..., K\}$ and all action dimension $n \in \{1, ..., N\}$ in parallel. This can be seen as extending the idea of Seyde et al. (2023), which learns decentralized Q-networks for action dimensions, into action sequence dimension. As we mentioned in Section 3.1, we find this simple scheme works well on challenging tasks with high-dimensional action spaces.

Architecture Let \mathbf{e}_k denote an one-hot encoding for k. For each level l, we construct features for each sequence step k as $\mathbf{h}_{t,k}^l = [\mathbf{h}_t, \mathbf{a}_{t+k-1}^{l-1}, \mathbf{e}_k]$. We encode each $\mathbf{h}_{t,k}^l$ with a shared MLP network and process them through GRU (Cho et al., 2014) to obtain $\mathbf{s}_{t,k}^l = f_{\theta}^{\mathsf{GRU}}(f_{\theta}^{\mathsf{MLP}}(\mathbf{h}_{t,1}^l), ..., f_{\theta}^{\mathsf{MLP}}(\mathbf{h}_{t,k}^l))$. We use a shared projection layer to map each $\mathbf{s}_{t,k}^l$ into Q-values at each sequence step k, *i.e.*, $\{Q_{\theta}^{l,k}(\mathbf{o}_t, a_{t+k-1}^{l,n}, \mathbf{a}_{t:t+K}^{l-1})\}_{n=1}^N = f_{\theta}^{\mathsf{proj}}(\mathbf{s}_{t,k}^l)$. We note that we compute Q-values for all dimensions $n \in \{1, ..., N\}$ at the same time with a big linear layer, which follows the design of Seo et al. (2024).