Contents

A Foundation of Diffusion Models   **2**

    A.1 SDEs/ODEs and Solvers 2

    A.2 Guided Sampling Methods 2

B Related Works   **3**

C Details of Experimental Setup   **3**

    C.1 Offline Reinforcement Learning Environments and Datasets 3

    C.2 Offline Imitation Learning Environments and Datasets 4

D Additional Experiments   **5**

    D.1 Impact of Model Size in RL Benchmarks 5

    D.2 Impact of Diffusion Backbones and Sampling Steps (Full Results) 6

    D.3 Additional Analyses of DMs in IL Benchmarks 7

E Experimental Details   **7**

    E.1 Computing Resources 7

    E.2 Evaluation Metircs 8

    E.3 Algorithm Hyperparameters 8

F Implemented Diffusion Models   **8**

    F.1 DDPM/DDIM/DPM-Solver/DPM-Solver++ 8

    F.2 EDM 10

    F.3 Rectified Flow 11

G Implemented Algorithms   **11**

    G.1 Diffusion Planners 11

    G.2 Diffusion Polices 12

    G.3 Diffusion Data Synthesizers. 13

H Limitations, Challenges, and Future Directions   **13**

I Potential Social Impact   **14**

J License   **14**

1

28 # A Foundation of Diffusion Models

29 ## A.1 SDEs/ODEs and Solvers

Assume a D-dimensional random variable $\boldsymbol{x}_0 \sim \mathbb{R}^D$ with an unknown distribution $q_0(\boldsymbol{x}_0)$ [1]. Diffusion Models (DMs) [24, 47] define a *forward process* $\{\boldsymbol{x}_t\}_{t \in [0,T]}$ with $T > 0$ by the *noise schedule* $\{\alpha_t, \sigma_t\}_{t \in [0,T]}$, such that $\forall t \in [0, T]$, $\boldsymbol{x}_t$ satisfies

$$\boldsymbol{x}_t = \alpha_t \boldsymbol{x}_0 + \sigma_t \boldsymbol{\epsilon}, \ \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I}), \tag{1}$$

where $\alpha_t, \sigma_t \in \mathbb{R}^+$ are differentiable functions of t and the *signal-to-noise-ratio* (SNR) α_t^2/σ_t^2 is strictly decreasing w.r.t t. The forward process in Equation (1) can also be described as a stochastic differential equation (SDE) for any $t \in [0, T]$ [24]:

$$\mathrm{d}\boldsymbol{x}_t = f(t)\boldsymbol{x}_t \mathrm{d}t + g(t)\mathrm{d}\boldsymbol{w}_t, \ \boldsymbol{x}_0 \sim q_0(\boldsymbol{x}_0), \tag{2}$$

where $\boldsymbol{w}_t \in \mathbb{R}^D$ is the standard Wiener process, and $f(t) = \frac{\mathrm{d}\log\alpha_t}{\mathrm{d}t}, g^2(t) = \frac{\mathrm{d}\sigma_t^2}{\mathrm{d}t} - 2\sigma_t^2 \frac{\mathrm{d}\log\alpha_t}{\mathrm{d}t}$. The SDE forward process in Equation (2) has an equivalent *reverse process* from time T to 0 [47]:

$$\mathrm{d}\boldsymbol{x}_t = [f(t)\boldsymbol{x}_t - g^2(t)\nabla_{\boldsymbol{x}}\log q_t(\boldsymbol{x}_t)]\mathrm{d}t + g(t)\mathrm{d}\bar{\boldsymbol{w}}_t, \ \boldsymbol{x}_T \sim q_T(\boldsymbol{x}_T), \tag{3}$$

where $\bar{w}_t$ is a standard Wiener process in the reverse time. One can sample $q_0(x_0)$ by directly solving the SDE in **??**, in which the only unknown term is the *score function* $\nabla_{\boldsymbol{x}}\log q_t(\boldsymbol{x}_t)$. In practice, a neural network $\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t)$ parameterized by θ can be trained to approximate the scaled score function $-\sigma_t \nabla_{\boldsymbol{x}}\log q_t(\boldsymbol{x}_t)$ by minimizing the score matching loss [19, 46, 47]:

$$\mathcal{L}(\theta) := \mathbb{E}_{t \sim \mathrm{Uniform}(0,T), \boldsymbol{x}_t \sim q_t(\boldsymbol{x}_t)} \left[\|\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t) + \sigma_t \nabla_{\boldsymbol{x}}\log q_t(\boldsymbol{x}_t)\|_2^2 \right] \tag{4}$$

$$= \mathbb{E}_{t \sim \mathrm{Uniform}(0,T), \boldsymbol{x}_0 \sim q_0(\boldsymbol{x}_0), \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})} \left[\|\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t) - \boldsymbol{\epsilon}\|_2^2 \right]. \tag{5}$$

Since $\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t)$ can be considered as a predicted Gaussian noise added to $\boldsymbol{x}_t$, it is usually called the *noise prediction model*. With a well-trained noise prediction model, SDE in **??** can be solved using numerical solvers, and DDPM [19] is one such method. However, numerical solvers require discretization from T to 0, in which the randomness of the Wiener process limits the step size [25]. For faster sampling, one can solve the following *probability flow* ODE, which is proven to have the same marginal distribution as that of the SDE for any $t \in [0, T]$ [47]:

$$\frac{\mathrm{d}x_t}{\mathrm{d}t} = f(t)\boldsymbol{x}_t - \frac{1}{2}g^2(t)\nabla_{\boldsymbol{x}}\log q_t(\boldsymbol{x}_t), \ \boldsymbol{x}_T \sim q_T(\boldsymbol{x}_T). \tag{6}$$

DDIM [46] discretizes the ODE to the first order for solving, achieving almost no loss in quality with fewer sampling steps. DPM-Solver [34, 35] leverages the semi-linearity of diffusion ODEs in **??** for exact solutions, eliminating errors in the linear terms, resulting in a higher sample quality. Some works also reformulate the framework. EDM [23] optimizes the design choices from a perspective of noise schedule and uses a specially designed score function preconditioning to improve the sample quality. Rectified flow [33], on the other hand, designs a straight probability flow ODE from the optimal transport (OT) perspective, which can straighten itself through *reflow* procedure. The straight property of Rectified flow allows high-quality generation in very few sampling steps.

A.2 Guided Sampling Methods

Guided sampling methods aim to draw samples from $q_0(x_0|y)$ to generate outputs with the characteristics of the label y. Depending on whether an additional classifier needs to be trained, guided sampling methods are divided into two categories: classifier guidance (CG) [5] and classifier-free guidance (CFG) [20].

Classifier Guidance: For conditional sampling, the score function needs to be changed to $\nabla_{\boldsymbol{x}}\log q_t(\boldsymbol{x}_t|\boldsymbol{y})$, which can be decomposed with the Bayes Theorem:

$$\nabla_{\boldsymbol{x}}\log q_t(\boldsymbol{x}_t|\boldsymbol{y}) = \nabla_{\boldsymbol{x}}\log q_t(\boldsymbol{x}_t) + \nabla_{\boldsymbol{x}}\log q_t(\boldsymbol{y}|\boldsymbol{x}_t), \tag{7}$$

---

[1]To ensure clarity, we establish the convention that the subscript t denotes the timestep in the diffusion process, while the superscript t represents the timestep in sequential decision-making problem.

2

where the first term can be approximated by the noise prediction model, and the second term is a *noising classifier* that predicts the label y of the corrupt data $\boldsymbol{x}_t$. In practice, an additional neural network $\mathcal{C}_\phi(\boldsymbol{x}_t, t, \boldsymbol{y})$ is trained to approximate $\log q_t(\boldsymbol{y}|\boldsymbol{x}_t)$, and its gradient is computed to guide sampling process:

$$\bar{\boldsymbol{\epsilon}}_\theta(\boldsymbol{x}_t, t, \boldsymbol{y}) = \boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t) - w\sigma_t\nabla_{\boldsymbol{x}}\mathcal{C}_\phi(\boldsymbol{x}_t, t, \boldsymbol{y}), \tag{8}$$

where w stands for the guidance scale. A larger value of w sharpens the classifier, amplifying the influence of the label y.

Classifier-free Guidance: According to Equation (7), the gradient of the classifier $\nabla_{\boldsymbol{x}}\log q_t(\boldsymbol{y}|\boldsymbol{x}_t)$ can be written to $\nabla_{\boldsymbol{x}}\log q_t(\boldsymbol{x}_t|\boldsymbol{y}) - \nabla_{\boldsymbol{x}}\log q_t(\boldsymbol{x}_t)$. By training a conditional noise prediction model $\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t, \boldsymbol{y})$, the sampling process can be guided with no additional classifier:

$$\bar{\boldsymbol{\epsilon}}_\theta(\boldsymbol{x}_t, t, \boldsymbol{y}) = \boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t) - w\sigma_t\nabla_{\boldsymbol{x}}\log q_t(\boldsymbol{y}|\boldsymbol{x}_t) = \boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t) + w(\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t, \boldsymbol{y}) - \boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t)) \tag{9}$$

where $\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t) = \boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t, \Phi)$ is approximated by the noise prediction model conditioned on a pre-specified label Φ standing for non-conditioning. Although CFG can generate trajectories specific to condition $\boldsymbol{y}$, it may cause the agent to reject higher likelihood trajectories in sequential environments, resulting in a performance drop [40]. Therefore, some methods [4, 40, 50] set the guidance weight w to 1, i.e., no guidance paradigm.

B Related Works

In recent years, DMs have demonstrated promising performance in various domains [51, 45, 32, 24], giving rise to several high-quality DM libraries, such as Diffusers [49] and Stable Diffusion [44]. These open-source libraries have significantly promoted research and applications in related fields. However, unfortunately, these libraries are designed for multimedia such as image, audio, and video generation, lacking adaptation for decision-making tasks. This is likely because DMs play diverse roles in decision-making, with various usage patterns and many unique mechanism incorporations, creating a gap in the multimedia generation paradigm. A library specially designed for decision-making is currently missing, and most research codebases are inherited from a few pioneering studies [21, 50, 4]. While effective, their algorithm-specific mechanisms and tightly coupled system architecture make it challenging for customized development.

`CleanDiffuser` aims to provide an "easy-to-hack" starter kit for research needs, offering researchers more exploration possibilities. We draw from the experience of many open-source decision-making libraries. For example, we emulate stable-baselines3 [43] to carefully reproduce results to provide practitioners with reliable baselines for method comparison. However, we inject more modular design to encourage users to freely design and modify. We also follow CORL [48] in designing clean and logically clear pipelines for readability, but, considering the complexity of DMs, abandon the one-file-from-scratch approach and opt for a one-file pipeline approach to offer rich examples of how to utilize `CleanDiffuser` building blocks to implement decision-making algorithms. Additionally, we follow Ray [30] in providing ample parameter selection interfaces within modules, making it easy for users unfamiliar with the internal implementation to customize effortlessly. In summary, `CleanDiffuser` is not only the first open-sourced modularized DM library tailored for decision-making algorithms but also a new library that draws on the advanced experiences of many open-source decision-making libraries.

C Details of Experimental Setup

C.1 Offline Reinforcement Learning Environments and Datasets

We evaluate 7 diffusion-based RL algorithms implemented with CleanDiffuser on 15 offline RL tasks from 3 benchmarks, including locomotion, manipulation, and navigation. These tasks are widely recognized and extensively used in offline RL settings [27, 10, 26, 16, 50, 22, 21, 1, 7, 29, 18], enjoying significant acceptance within the research community. Visualization of these tasks is

3
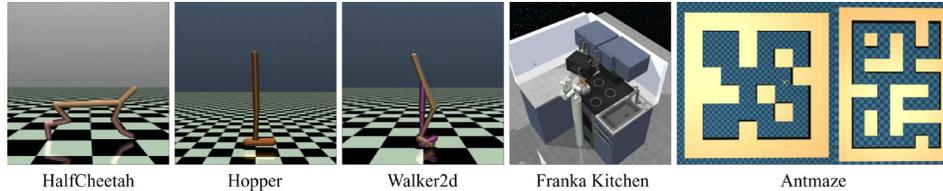
Figure 1: **Visualization of Offline Reinforcement Learning Environments.**

presented in Figure 1. These tasks come from the D4RL benchmark, in which the datasets are licensed under the Creative Commons Attribution 4.0 License (CC BY), and the code is licensed under the Apache 2.0 License.

Gym-MuJoCo [2] consists of three popular offline RL locomotion tasks (HalfCheetah, Hopper, Walker2d), which require controlling three Mujoco robots to achieve maximum movement speed while minimizing energy consumption under stable conditions. D4RL [9] benchmark provides three different quality levels of offline datasets: "medium" containing demonstrations of medium-level performance; "medium-replay" containing all recordings in the replay buffer observed during training until the policy reaches "medium" performance; and "medium-expert" which combines "medium" and "expert" level performance equally.

Franka Kitchen [11] requires controlling a realistic 9-DoF Franka robot arm to complete several household tasks in a kitchen environment. Algorithms are trained on "partial" and "mixed" datasets. The "partial" and "mixed" datasets consist of undirected data, where the robot performs subtasks that are not necessarily related to the goal configuration. In the "partial" dataset, a subset of the dataset is guaranteed to solve the task, meaning an imitation learning agent may learn by selectively choosing the right subsets of the data. The "mixed" dataset contains no trajectories that solve the task completely, and the RL agent must learn to assemble the relevant sub-trajectories. This dataset requires the highest degree of generalization in order to succeed.

Antmaze [9] requires controlling the 8-DoF "Ant" quadruped robot to complete maze navigation tasks. In the offline dataset, the robot only receives a reward upon reaching the goal, and the dataset contains many trajectory segments that do not lead to the endpoint, making it a difficult decision task with sparse rewards and a long horizon. The success rate of reaching the endpoint is used as the evaluation score, and common offline RL algorithms often struggle to achieve good performance.

C.2 Offline Imitation Learning Environments and Datasets
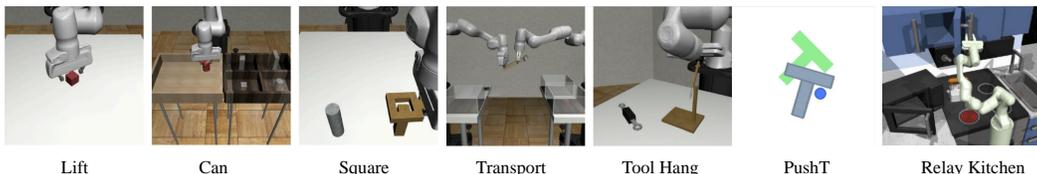
Figure 2: **Visualization of Offline Imitation Learning Environments.**

We evaluate 2 diffusion-based IL algorithms implemented with CleanDiffuser on 22 imitation learning tasks from 4 benchmarks, with both state and image-based observation inputs. Among them, Relay Kitchen and Robomimic support both velocity and position control. Each algorithm is trained with its best-performing action space. We provide task summary in Table 1, visualization in Figure 2, and more details below:

PushT [8] requires pushing a T-shaped block (gray) to a fixed target (red) with a circular end-effector. The task requires exploiting complex and contact-rich object dynamics to push the T block precisely, using point contacts. In this paper, we used three variants. "PushT" env has a five-dimensional state space, including the proprioception for end-effector location (`agent_x, agent_y`) and the

4

xy coordinates and angles of the blocks (`block_x`, `block_y`, `block_angle`). "PushT-keypoints" env includes nine 2D key points obtained from the T-block's ground truth attitude and proprioception for end-effector location. "Pusht-image" env observes the end-effector location and the top view of the RGB image. This benchmark is licensed under the Apache-2.0 License.

Relay Kitchen is proposed in Relay Policy Learning [11], commonly used to evaluate imitative learning ability. The environment consists of a 9 DoF position-controlled Franka robot interacting with a kitchen scene that includes an openable microwave, four turnable oven burners, an oven light switch, a freely movable kettle, two hinged cabinets, and a sliding cabinet door. The "relay" dataset contains 566 human demonstrations, each completing four tasks in arbitrary order. The goal is to execute as many tasks as possible, regardless of order, showcasing both short-horizon and long-horizon multimodality. This benchmark is licensed under the Apache-2.0 License.

Robomimic [37] requires controlling a robot arm to complete complex manipulation tasks from a few human demonstrations. Due to the non-Markovian nature of human demonstrations and the demonstration quality variance, learning from human datasets is significantly more challenging than learning from machine-generated datasets. Proficient-Human (PH) and Multi-Human (MH) datasets are collected by humans through remote teleoperation. The PH datasets consist of 200 demonstrations collected by a single, experienced teleoperator, while the MH datasets consist of 300 demonstrations collected by 6 teleoperators of varying proficiency, each of which provided 50 demonstrations. The benchmark consists of 5 PH tasks (Lift, Can, Square, Tool_hang, Transport) and 4 MH tasks (Lift, Can, Square, Transport). Each task has both state and image-based observation inputs. This benchmark is licensed under the MIT License.

To the best of our knowledge, the datasets and benchmarks we have used do not contain personally identifiable information or offensive content in both previous works and our works.

Table 1: **Imitation Learning Task Summary.** Obs Shape represents the low dimensional state space dimension; Image Shape represents the observation resolution of multi-view images (Camera views x W x H). PH: proficient-human demonstration, MH: multi-human demonstration, Steps: max episode steps.

Task	Low Dim Tasks	Image Tasks		Action Dim	PH Demonstration	MH Demonstration	Max Steps
	Obs Shape	Obs Shape	Image Shape				
PushT	5	N/A	N/A	2	200	N/A	300
PushT-Keypoint	20	N/A	N/A	2	200	N/A	300
PushT-Image	N/A	2	1x96x96	2	200	N/A	300
Relay Kitchen	60	N/A	N/A	9	656	N/A	280
Lift	19	9	2x84x84	7	200	300	400
Can	23	9	2x84x84	7	200	300	400
Square	23	9	2x84x84	7	200	300	500
Transport	59	18	4x84x84	7	200	300	700
Tool_hang	53	9	2x240x240	7	200	N/A	700

D Additional Experiments

D.1 Impact of Model Size in RL Benchmarks

Table 2: **Impact of Model Size in RL Benchmarks.** Performance of DD and IDQL with varying model sizes. Results correspond to the mean and standard error over 150 episode seeds.

Environment	DD			IDQL		
Model Size	4M	15M	60M	1.6M	6M	25M
HalfCheetah-m	45.3 ± 0.3	44.5 ± 0.1	$\mathbf{47.1 \pm 0.1}$	51.5 ± 0.1	51.5 ± 0.1	$\mathbf{51.7 \pm 0.1}$
Kitchen-m	56.5 ± 5.8	$\mathbf{80.5 \pm 4.1}$	27.7 ± 2.1	66.5 ± 4.1	$\mathbf{69.2 \pm 1.0}$	67.5 ± 1.8
Antmaze (mp for DD, lp for IDQL)	8.0 ± 4.3	$\mathbf{26.0 \pm 5.9}$	22.7 ± 6.6	48.7 ± 4.7	52.0 ± 5.7	$\mathbf{54.0 \pm 4.3}$

There is a significant disparity in network model sizes used by diffusion-based decision-making algorithms. For instance, the official implementation of DD utilizes around 60M parameters [1], while Diffuser uses 4M [21], and IDQL [16] has approximately only 1.6M parameters. These works have

5

limited discussion on the impact of model size. Therefore, we aim to explore the approximate scale of parameter sizes required for diffusion-based decision-making algorithms to function effectively. In this experiment, we test DD and IDQL at three different model sizes, starting from the default parameter size used in the main experiments and gradually increasing the parameter size by four times. The performance of the algorithms is evaluated on three tasks including locomotion, manipulation, and navigation. Results are presented in Table 2. We find that, apart from the performance of DD on Kitchen-m and Antmaze-mp, increasing the model size does not lead to significant performance gains in other cases. However, even with the performance gains brought by model size, DD can not entirely catch up with the performance of IDQL, indicating that the dominant effect on performance is still primarily driven by the algorithm rather than the model size.

D.2   Impact of Diffusion Backbones and Sampling Steps (Full Results)
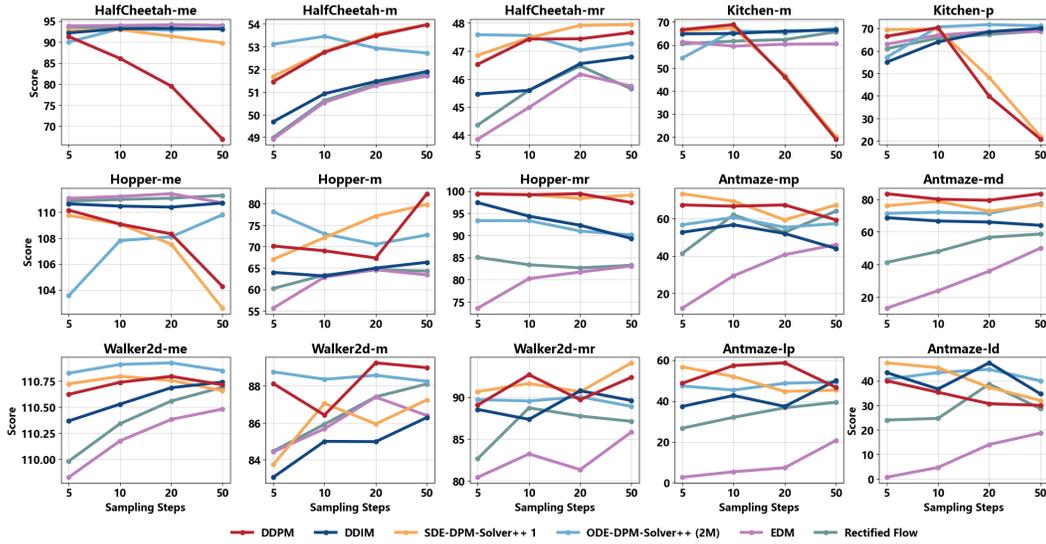
Figure 3: **Full D4RL Results of IDQL.** Performance of IDQL with various diffusion backbones and varying sampling steps. Results correspond to the mean over 150 episode seeds.
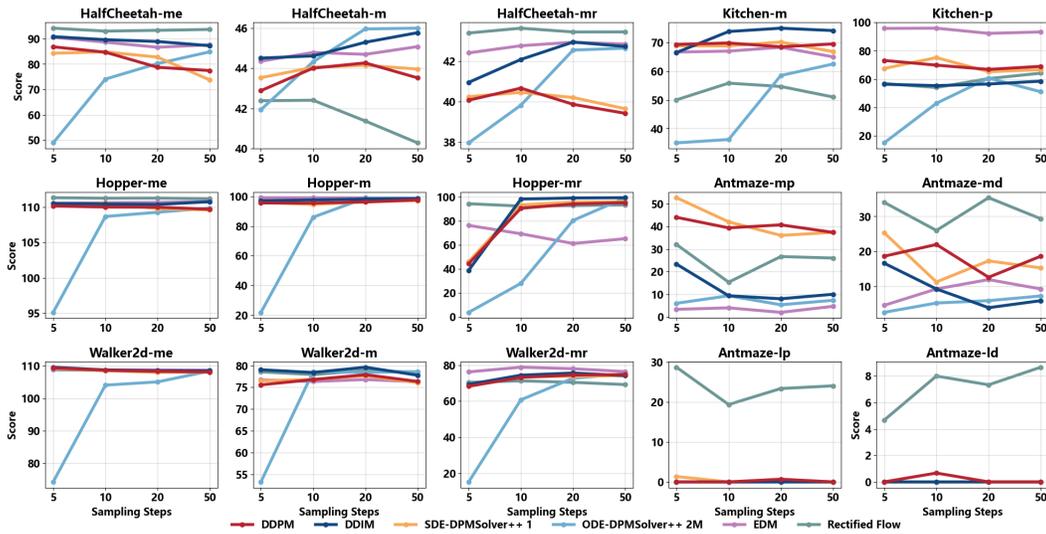
Figure 4: **Full D4RL Results of DD.** Performance of DD with various diffusion backbones and varying sampling steps. Results correspond to the mean over 150 episode seeds.

6

179 Due to space limitations in the main text, we present the full results of IDQL and DD on D4RL in
180 Figure 3 and Figure 4. The algorithms are trained for 1×10^6 gradient steps, and the sampling steps for
181 DD are set to 5, with other hyperparameters consistent with default settings. This experiment selects
182 DDPM, DDIM, SDE-DPM-Solver++ 1, ODE-DPM-Solver++ (2M), EDM, and Rectified Flow as the
183 diffusion/solver backbones. We select DDPM and DDIM because they are the first-order discretization
184 of diffusion reverse SDE/ODE, respectively [47, 46]. We do not choose DPM-Solver because its
185 first-order solver is equivalent to DDIM [34], and higher-order solvers may cause instability under
186 guidance [35]. For DPM-Solver++, we select a first-order SDE solver, SDE-DPM-Solver++ 1,
187 and a second-order ODE solver, ODE-DPM-Solver++ (2M). Since higher-order solvers can lead
188 to instability, they are therefore not chosen. We select EDM and Rectified Flow because they have
189 achieved excellent results in image generation but have not been widely used in the decision-making
190 domain, to the best of our knowledge. Thanks to `CleanDiffuser`'s support for various solvers
191 and varying sampling steps, the results for DDPM, DDIM, SDE-DPM-Solver++ 1, and ODE-DPM-
192 Solver++ (2M) only require training one single model. Additionally, using different sampling steps
193 does not require additional training. These features provide a great convenience for conducting
194 ablation experiments. We believe these features of `CleanDiffuser` can also benefit future research
195 efforts.

D.3 Additional Analyses of DMs in IL Benchmarks

197 Using the low-dim lift-ph task with 50 sample
198 steps in Robomimic as a reference, we present
199 the number of parameters and inference time for
200 each variant of DiffusionPolicy, DiffusionBC,
201 and ACT in table 3. Although `Chi_UNet1d` ex-
202 hibits the best performance in many IL tasks,
203 it has the largest model size and the slowest
204 inference speed. Larger model size results in
205 higher training costs, and in many real-world
206 applications that require real-time inference, we
207 need to make trade-offs between inference speed
208 and performance. Compared to the transformer-
209 based ACT algorithm, all structures of the dif-
210 fusion policy exhibit slower sampling speeds
211 because the denoising process requires multiple
212 forwards for neural networks. This is also an

Table 3: **The Model Size and Inference Time of Dif-fusionPolicy and DiffusionBC in Low-Dim Lift-ph.** DiffusionPolicy uses 50 sampling steps across the experiments, and DiffusionBC incorporates 8 additional Diffusion-X sampling steps.

Algorithm	Model Size (M)	Inference Time (s)
DiffusionPolicy w/ `Chi_UNet1d`	68.91	0.405
DiffusionPolicy w/ `Chi_TFM`	9.50	0.343
DiffusionPolicy w/ `DiT1d`	16.59	0.194
DiffusionBC w/ `DiT1d`	16.59	0.217
DiffusionBC w/ `Pearce_MLP`	0.83	0.062
ACT	7.83	0.006

213 important challenge that limits the application of DMs for decision-making. We also note that
214 DiffusionBC is slower than DiffusionPolicy when using the same network architecture and model
215 size, as DiffusionBC performs 8 additional steps of Diffusion-X sampling to mitigate OOD issues.
216 Although the best-performing `Chi_UNet1d` model uses a considerable model size, simply increasing
217 the Transformer-based DMs like `DiT1d` can sometimes harm performance. We discuss this in detail
218 in appendix D.1, which is also consistent with the experimental observations of the [4]. Finding the
219 optimal model size in applications remains an open research question.

E Experimental Details

E.1 Computing Resources

222 RL experiments are conducted on a server equipped with 2 Intel(R) Xeon(R) Gold 6326 CPUs @
223 2.90GHz and 8 NVIDIA GeForce RTX3090 GPUs, and a server equipped with 2 Intel(R) Xeon(R)
224 Gold 6326 CPUs @ 2.90GHz and 8 NVIDIA GeForce RTX2080Ti GPUs. IL experiments are
225 conducted on a server equipped with 2 Intel(R) Xeon(R) Gold 6338 CPUs @ 2.00GHz and 8 NVIDIA
226 A800 GPUs, and a server equipped with 2 Intel(R) Xeon(R) Gold 6338 CPUs @ 2.00GHz and 4
227 NVIDIA GeForce RTX3090 GPUs.

7

E.2 Evaluation Metircs

In the D4RL benchmark, the scores are normalized to the range between 0 and 100 with expert-normalized scores $= 100 \times \frac{\text{score} \times \text{random\_score}}{\text{expert\_score-random\_score}}$ [9]. As for IL benchmarks, we report target area coverage as scores in the PushT benchmark and success rate in the Robomimic benchmark. In the Relay Kitchen environment, since the vast majority of human demonstrations can only complete 4 subtasks, we denote the success rate of completing the i-th subtask as p_i and report the average success rate as score $= (p_1 + p_2 + p_3 + p_4)/4$.

E.3 Algorithm Hyperparameters

Unless stated otherwise, we utilize default hyperparameters from the official implementations for most algorithms and datasets. **Configuration files and hyperparameters for each algorithm and environment are available in YAML format on our GitHub repository for reproducibility.**

Key hyperparameters for each offline RL algorithm are presented in Table 4, and each offline IL algorithm in Table 5. We also reproduce the Transformer-based ACT [52] algorithm based on the official implementation, the key hyperparameters are in Table 50000.

Table 4: **Hyperparameters for Diffusion Planners, Diffusion Policies and Diffusion Data Synthesizer for RL.**

Hyperparameter	Diffuser	DD	AdaptDiffuser	DQL	EDP	IDQL	SynthER
Architecture	Janner_UNet	DiT	Janner_UNet	DQL_MLP	DQL_MLP	LNResnet	LNResnet
Diffusion Model	DDPM	DDIM	DDPM	DDPM	DPM-Solver++ (2M)	DDPM	DDIM
Sampling Steps	20	20	20	5	15	5	128
Horizon	64 (Antmaze)	64 (Antmaze)	64 (Antmaze)	1	1	1	1
	32 (Otherwise)	32 (Otherwise)	32 (Otherwise)				
Temperature	0.5	0.5	0.5	0.5	0.5	0.5	1.0
Gradient Steps	1e6	1e6	1e6	2e6	2e6	2e6	1e5
Batch Size	64	64	64	256	256	256	256
Learning Rate	3e-4	3e-4	3e-4	3e-4	3e-4	3e-4	3e-4
N candidates	64	1	64	50	50	256	N/A

Table 5: **Hyperparameters for DiffusionPolicy and DiffusionBC in Low-Dim and Image Tasks.**

Hyperparameters	DiffusionPolicy			DiffusionBC	
Architecture	Chi_UNet1d	Chi_Transformer	DiT1d	Pearce_MLP	DiT
Diffusion Model	DDPM	DDPM	DDPM	DDPM	DDPM
Sampling Steps	5 (PushT)	5 (PushT)	5 (PushT)	50	50
	50 (Otherwise)	50 (Otherwise)	50 (Otherwise)		
Horizon	16	10	10	2	2
Obs Steps	2	2	2	2	2
Action Steps	8	8	8	1	1
Gradient Steps	1e6	1e6	1e6	1e6	1e6
Batch Size	256 (Low dim)	256 (Low dim)	256 (Low dim)	512 (Low dim)	512 (Low dim)
	64 (Image)	64 (Image)	64 (Image)	64 (Image)	64 (Image)
Temperature	1.0	1.0	1.0	1.0	1.0
Learning Rate	1e-4	1e-4	1e-4	1e-3	5e-4
Extra Sample Steps	N/A	N/A	N/A	8	8
Control Mode	Pos	Pos	Pos	Vel	Vel

F Implemented Diffusion Models

F.1 DDPM/DDIM/DPM-Solver/DPM-Solver++

Applying Solvers with One Score Function. Due to the generation processes of DDPM [19], DDIM [46], DPM-Solver [34], and DPM-Solver++ [35] can all be expressed using the same diffusion SDE/ODE [47], utilizing the same noise schedule, training just one noise predictor model enables the use of these four solvers for sampling. Recall that the diffusion ODE with noise prediction model is:

$$\frac{\mathrm{d}x_t}{\mathrm{d}t} = f(t)\boldsymbol{x}_t + \frac{g^2(t)}{2\sigma_t}\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t). \tag{10}$$

8

Table 6: **Hyperparameters for ACT in Low-Dim and Image Tasks.**

Hyperparameters	Value
Learning Rate	1e-5
Batch Size	256 (Low dim) / 64 (Image)
# Encoder Layers	4
# Decoder Layers	7
Feedforward Dimension	256
Hidden Dimension	256
# Heads	8
Chunk size	16
Beta	10
Gradient Steps	1e6
Control Mode	Vel (Kitchen) / Pos (Otherwise)

Substituting $f(t) = \frac{\mathrm{d}\log\alpha_t}{\mathrm{d}t}, g^2(t) = \frac{\mathrm{d}\sigma_t^2}{\mathrm{d}t} - 2\sigma_t^2\frac{\mathrm{d}\log\alpha_t}{\mathrm{d}t}$, and conducting first-order discretization result in a recursive formula:

$$\boldsymbol{x}_t - \boldsymbol{x}_s = \frac{\alpha_t - \alpha_s}{\alpha_s}\boldsymbol{x}_s + \frac{1}{2\sigma_s}\left[2\sigma_s(\sigma_t - \sigma_s) - 2\frac{\sigma_s^2}{\alpha_s}(\alpha_t - \alpha_s)\right]\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_s) \tag{11}$$

$$\boldsymbol{x}_t = \frac{\alpha_t}{\alpha_s}\boldsymbol{x}_s - \alpha_t\left(\frac{\sigma_s}{\alpha_s} - \frac{\sigma_t}{\alpha_t}\right)\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_s, s) \tag{12}$$

$$\boldsymbol{x}_t = \alpha_t\left(\frac{\boldsymbol{x}_t - \sigma_t\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_s, s)}{\alpha_s}\right) + \sqrt{\sigma_s^2}\epsilon_\theta(\boldsymbol{x}_s, s), \tag{13}$$

where t and s are the next and current sampling steps. Equation (13) is DDIM update [46]. By introduce $\beta_s = (\sigma_t/\sigma_s)\sqrt{1 - \alpha_s^2/\alpha_t^2}$, the generative process of DDPM is:

$$\boldsymbol{x}_t = \alpha_t\left(\frac{\boldsymbol{x}_t - \sigma_t\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_s, s)}{\alpha_s}\right) + \sqrt{\sigma_s^2 - \beta_s^2}\epsilon_\theta(\boldsymbol{x}_s, s) + \beta_s\boldsymbol{\epsilon}_s, \tag{14}$$

where $\boldsymbol{\epsilon}_s \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$ is standard Gaussian noise independent of $\boldsymbol{x}_s$. DPM-Solver leverages the semi-linearity of the diffusion ODE and formulates the exact solution by the "variation of constants" formula:

$$\boldsymbol{x}_t = e^{\int_s^t f(\tau)\mathrm{d}\tau}\boldsymbol{x}_s + \int_s^t\left(e^{\int_\tau^t f(r)\mathrm{d}r}\frac{g^2(\tau)}{2\sigma_\tau}\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_\tau, \tau)\right)\mathrm{d}\tau \tag{15}$$

$$\boldsymbol{x}_t = \frac{\alpha_t}{\alpha_s}\boldsymbol{x}_s - \alpha_t\int_s^t\frac{\mathrm{d}\lambda_\tau}{\mathrm{d}\tau}\frac{\sigma_\tau}{\alpha_\tau}\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_\tau, \tau)\mathrm{d}\tau, \tag{16}$$

where $\lambda_t := \log(\alpha_t/\sigma_t)$ is the log-*signal-to-noise-ratio* (log-SNR). This formulation eliminates the approximation error of the linear term since it is exactly computed, and the non-linear term can be approximated using its Talor expansion:

$$\boldsymbol{x}_t = \frac{\alpha_t}{\alpha_s}\boldsymbol{x}_s - \alpha_t\sum_{n=0}^{k-1}\boldsymbol{\epsilon}_\theta^{(n)}(\boldsymbol{x}, s)\int_{\lambda_s}^{\lambda_t}e^{-\lambda}\frac{(\lambda - \lambda_s)^n}{n!}\mathrm{d}\lambda + \mathcal{O}((\lambda_t - \lambda_s)^{k+1}). \tag{17}$$

In `CleanDiffuser`, we have implemented only *DPM-Solver-1*, corresponding to the k=1 scenario in Equation (17), as guided sampling tends to make high-order solvers unstable [35], leading to poor performance in decision-making tasks. DPM-Solver++ alleviates this instability issue by using a data prediction model $\boldsymbol{x}_\theta(\boldsymbol{x}_t, t)$ instead of the noise prediction model $\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t)$, transforming the generative process into:

$$\boldsymbol{x}_t = \frac{\sigma_t}{\sigma_s}\boldsymbol{x}_s + \sigma_t\sum_{n=0}^{k-1}\boldsymbol{x}_\theta^{(n)}(\boldsymbol{x}, s)\int_{\lambda_s}^{\lambda_t}e^{\lambda}\frac{(\lambda - \lambda_s)^n}{n!}\mathrm{d}\lambda + \mathcal{O}((\lambda_t - \lambda_s)^{k+1}), \tag{18}$$

where $\boldsymbol{x}_\theta(\boldsymbol{x}_t, t)$ is trained to predict the original data $\boldsymbol{x}_0$ from the perturbed data $\boldsymbol{x}_s$. In `CleanDiffuser`, we have implemented DPM-Solver++ for $k \leq 2$, as it already yields satisfactory results at $k = 2$, while higher-order solvers may still lead to instability.

9

Although the data prediction model can mitigate the instability issue caused by guided sampling and easily clip data to address the "train-test mismatch" problem [35], there is still no definitive evidence in practice to determine the superiority of either the data prediction model or the noise prediction model. In CleanDiffuser, we provide users with the option to choose between these two prediction models and use the approximation $\boldsymbol{x}_t \approx \alpha_t \boldsymbol{x}_\theta(\boldsymbol{x}_t, t) + \sigma \boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t)$ to seamlessly switch between the two formulations to cater to the requirements of different solvers.

Noise Schedules. CleanDiffuser provides two popular noise schedules by default: *Linear Noise Schedule* [19] and *Cosine Noise Schedule* [38]. The former defines:

$$\alpha_t = \exp\left(-\frac{(\beta_1 - \beta_0)}{4}t^2 - \frac{\beta_0}{2}t\right), \tag{19}$$

where $\beta_0 = 0.1$, $\beta_1 = 20$ and $\sigma_t = \sqrt{1 - \alpha_t^2}$. The diffusion SDE/ODE is solved between $[\epsilon, T]$, where $\epsilon = 0.001$ and $T = 1$ for numerical stability. The later schedule defines:

$$\alpha_t = \frac{\cos\left(\frac{\pi}{2} \cdot \frac{t+s}{1+s}\right)}{\cos\left(\frac{\pi}{2} \cdot \frac{s}{1+s}\right)} \tag{20}$$

where $s = 0.008$ and $\sigma_t = \sqrt{1 - \alpha_t^2}$. The diffusion SDE/ODE is solved between $[\epsilon, T]$, where $\epsilon = 0.001$ and $T = 0.9946$ for numerical stability. Beyond the two schedules, CleanDiffuser allows users to fully customize new noise schedules according to the specified format to explore algorithm performance.

F.2 EDM

EDM [23] rewrites the diffusion forward process in Equation (1) as:

$$\boldsymbol{x}_t = s_t(\boldsymbol{x}_0 + \sigma_t \boldsymbol{\epsilon}_t), \tag{21}$$

which can be interpreted as adding noise to a scaled version of the original data. By setting the scale $s_t \equiv 1$ to a constant, EDM obtains the following reverse process:

$$\frac{\mathrm{d}\boldsymbol{x}_t}{\mathrm{d}t} = -\dot{\sigma}_t \sigma_t \nabla_{\boldsymbol{x}} \log p(\boldsymbol{x}; \sigma_t)\mathrm{d}t, \tag{22}$$

where $p(\boldsymbol{x}; \sigma_t) = p_t(\boldsymbol{x})$. A data prediction model $D_\theta(\boldsymbol{x}; \sigma)$ is trained to approximate $\boldsymbol{x} + \sigma^2 \nabla_{\boldsymbol{x}} \log p(\boldsymbol{x}; \sigma)$ and results in a practical generative process:

$$\boldsymbol{x}_t = \boldsymbol{x}_s + (t - s) \cdot \left(\frac{\dot{\sigma}_s}{\sigma_s}\boldsymbol{x}_s - \frac{\dot{\sigma}_s}{\sigma_s}D_\theta(\boldsymbol{x}_s; \sigma_s)\right). \tag{23}$$

One feature of EDM is that it applies preconditioning to D_θ:

$$D_\theta(\boldsymbol{x}; \sigma) = c_{\text{skip}}(\sigma)\boldsymbol{x} + c_{\text{out}}(\sigma)F_\theta(c_{\text{in}}(\sigma)\boldsymbol{x}; c_{\text{noise}}(\sigma)). \tag{24}$$

where F_θ is the neural network to be trained, c_{skip} modulates the skip connection, c_{in} and c_{out} scale the input and output magnitudes, and c_{noise} maps noise level σ into a conditioning input for F_θ. F_θ is trained by minimizing the noising score matching loss:

$$\mathcal{L}(\theta; \sigma) = \mathbb{E}_{\boldsymbol{y} \sim p_{\text{data}}, \boldsymbol{n} \sim \mathcal{N}(\boldsymbol{0}, \sigma^2 \boldsymbol{I})} \left[\lambda(\sigma)\|D_\theta(\boldsymbol{y} + \boldsymbol{n}; \sigma) - \boldsymbol{y}\|_2^2\right], \tag{25}$$

where $\lambda(\sigma)$ is the loss weight. These coefficients are optimized to achieve the following objectives: (1) inputs of F_θ have unit variance, (2) training target of F_θ have unit variance, (3) c_{skip} can minimize c_{out} so that the errors of F_θ are amplified as little as possible, and (4) the loss of F_θ has a uniform weight across noise levels. The optimization results give the following design choices: $c_{\text{skip}} = \sigma_{\text{data}}^2/(\sigma^2 + \sigma_{\text{data}}^2)$, $c_{\text{out}} = \sigma \cdot \sigma_{\text{data}}/\sqrt{\sigma_{\text{data}}^2 + \sigma^2}$, $c_{\text{in}} = 1/\sqrt{\sigma_{\text{data}}^2 + \sigma^2}$, $c_{\text{noise}} = \log(\sigma)/4$, and $\lambda(\sigma) = (\sigma_{\text{data}}^2 + \sigma^2)/(\sigma_{\text{data}} \cdot \sigma)^2$.

Noise Schedule. CleanDiffuser provides only one default noise schedule, which is specially designed for EDM:

$$\sigma_t = t, \ t \in [\sigma_{\min}, \sigma_{\max}], \tag{26}$$

where $\sigma_{\min} = 0.002$ and $\sigma_{\max} = 80$.

10

F.3 Rectified Flow

Rectified flow [33] is an ODE on time $t \in [0, 1]$:

$$\frac{\mathrm{d}\boldsymbol{x}_t}{\mathrm{d}t} = \boldsymbol{v}_\theta(\boldsymbol{x}_t, t), \tag{27}$$

where the drift force $\boldsymbol{v}_\theta$ is trained to drive the flow to follow the direction $(\boldsymbol{x}_0 - \boldsymbol{x}_1)$ of the linear path pointing from $\boldsymbol{x}_1$ to $\boldsymbol{x}_0$ as much as possible, by solving a simple least squares regression problem:

$$\mathcal{L}(\theta) = \mathbb{E}_{\boldsymbol{x}_0 \sim p_0, \boldsymbol{x}_1 \sim p_1, t \sim \mathrm{Uniform}(0,1)} \left[\left\| (\boldsymbol{x}_0 - \boldsymbol{x}_1) - \boldsymbol{v}_\theta(\boldsymbol{x}_t, t) \right\|_2^2 \right], \tag{28}$$

where $\boldsymbol{x}_t = t\boldsymbol{x}_1 + (1-t)\boldsymbol{x}_0$. It achieves the mutual transformation of samples from two distributions p_0 and p_1, by solving Equation (26) forward or backward. Rectified flow possesses many favorable properties that allow it to continuously learn from its own sampled data to straighten the ODE flow, and this procedure is called *reflow*. The straighter the ODE flow, the fewer sampling steps are needed to achieve good generation quality. In an ideal scenario, if the flow becomes completely straight, then we have:

$$\boldsymbol{x}_t = t\boldsymbol{x}_1 + (1-t)\boldsymbol{x}_0 = \boldsymbol{x}_1 + (1-t)\boldsymbol{v}(\boldsymbol{x}_1, 1), \ \forall t \in [0,1], \tag{29}$$

which enables one-step sampling. The Rectified Flow implemented in `CleanDiffuser` has full functionality to transform samples from any two arbitrary probability distributions. By default, it follows the settings in diffusion models, where p_0 is the dataset distribution and p_1 is the standard Gaussian distribution.

G Implemented Algorithms

G.1 Diffusion Planners

Diffuser. [21] Diffuser is the first diffusion planning algorithm, and its paradigm has been widely adopted in subsequent diffusion planning algorithms. Diffuser generates state-action pair trajectories $\boldsymbol{x} = [x^\tau, \cdots, x^{\tau+H-1}]$ from:

$$p(\boldsymbol{x}|\mathcal{O}^{\tau:\mathcal{T}}) \propto p(\boldsymbol{x})p(\mathcal{O}^{\tau:\mathcal{T}}|\boldsymbol{x}) = p(\boldsymbol{x}) \prod_{t=\tau}^{\mathcal{T}} \exp(r(s^t, a^t)), \tag{30}$$

where $\mathcal{O}^{t_1:t_2}$ is a binary random variable denoting the optimality of a trajectory from t_1 to t_2, and $\mathcal{T}$ is the episode terminal time step of the trajectory [2]. Therefore, it is natural to define the classifier in CG as a reward function on perturbed trajectories:

$$\nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x}_t|\mathcal{O}^{\tau:\mathcal{T}}) = \nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x}_t) + \sum_{k=\tau}^{\mathcal{T}} \nabla_{s_t^k, a_t^k} r(s_t^k, a_t^k) = \nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x}_t) + \nabla_{\boldsymbol{x}} \mathcal{J}_\phi(\boldsymbol{x}_t, t), \tag{31}$$

where $\mathcal{J}_\phi(\boldsymbol{x}_t, t)$ is a neural network trained to predict the episodic cumulative reward $\sum_{k=\tau}^{\mathcal{T}} r(s_t^k, a_t^k)$ of the perturbed trajectory $\boldsymbol{x}_t$. At each inference step, given the current state s^k, Diffuser sets and freezes the first state of the trajectory as s^k and performs guided sampling in an inpainting manner to generate a set of trajectories $\{\boldsymbol{x}_0\}$. Subsequently, it identifies the optimal trajectory $\boldsymbol{x}_0^* = \arg\max_{\boldsymbol{x}_0} \mathcal{J}_\phi(\boldsymbol{x}_0, 0)$ that maximizes the episodic cumulative reward, and extracts the first action a^k in $\boldsymbol{x}_0^*$ to execute.

Decision Diffuser. [1] Decision Diffuser (DD) introduces another prominent framework that utilizes a state-only trajectory formulation and implements CFG by discarding the optimality variable $\mathcal{O}^{\tau:\mathcal{T}}$ in favor of directly employing normalized episodic cumulative reward $y = \sum_{t=\tau}^{\mathcal{T}} r(\boldsymbol{x})$ as the condition. As no additional reward predictor can be used for trajectory selection, DD generates only a single

---

[2] In previous works, authors typically consider only the trajectory cumulative reward as the generative condition, i.e. using $\mathcal{O}^{\tau:\tau+H-1}$, which overlooks future optimality. Their code implementations actually use the episodic cumulative reward, i.e. $\mathcal{O}^{\tau:\mathcal{T}}$. Therefore, we adopt this episodic cumulative reward expression.

11

trajectory at each inference step and employs an trained inverse dynamic model $\mathcal{I}_\phi$ to predict the action to be executed $a^t = \mathcal{I}_\phi(s^t, s^{t+1})$.

AdaptDiffuser. [31] Observing that the insufficient diversity of offline RL training data may limit the sample quality of DMs, AdaptDiffuser, an extension of Diffuser, proposes to utilize self-generated diverse synthetic expert data to fine-tune itself. The pipeline of AdaptDiffuser involves initially training a Diffuser as usual, then generating a large amount of synthetic expert data and using a discriminator to filter out high-quality data. Finally, fine-tuning is done on this dataset. This self-evolving process can be repeated multiple times to optimize the model, and different directions of model self-evolution can be controlled by designing different discriminators. The inference method of AdaptDiffuser is consistent with Diffuser, and its performance for seen tasks has been enhanced while also being able to adapt to unseen tasks.

G.2 Diffusion Polices

Diffusion Q-Learning. [50] Diffusion Q-learning (DQL) leverages the capability of DMs to model complex distributions, directly applying DDPM as the policy $\pi_\theta(a_0|s)$ in the RL actor-critic framework. Sampling from the policy is therefore equivalent to the denoising process of the diffusion model. The Bellman operator can be used to train the Q-value function of the diffusion policy:

$$\mathcal{L}(\phi) = \mathbb{E}_{(s^k, a^k, r, s^{k+1}) \sim \mathcal{D}, a_0^{k+1} \sim \pi_{\theta'}} \left[\left\| \left(r + \gamma \min_{i=1,2} Q_{\phi_i'}(s^{k+1}, a_0^{k+1}) \right) - Q_{\phi_i}(s^k, a^k) \right\|_2^2 \right], \quad (32)$$

where ϕ_1 and ϕ_2 represent the parameters of the double Q-learning trick, ϕ' and θ' represent the target networks. For policy optimization, DQL employs the most basic form of Offline RL optimization, which involves training the policy to maximize the Q-value while imitating behavior policies, using a weighting factor α to balance the influence of both aspects:

$$\mathcal{L}(\theta) = \mathcal{L}_{\text{score}}(\theta) - \alpha \cdot \mathbb{E}_{s \sim \mathcal{D}, a_0 \sim \pi_\theta} \left[Q_\phi(s, a_0) \right], \quad (33)$$

where $\mathcal{L}_{\text{score}}(\theta)$ is the score matching loss used for diffusion model training. As the scale of the Q-value function varies in different offline datasets, to normalize it, DQL sets $\alpha = \frac{\eta}{\mathbb{E}_{(s,a) \sim \mathcal{D}}[|Q_\phi(s,a)|]}$ and tunes η for loss term balance. The Q_ϕ in the denominator is only for normalization and not differentiated over.

Efficient Diffusion Policy. [22] Efficient Diffusion Policy (EDP) aims to address the significant computational overhead caused by iterative sampling and gradient computation during the training of the DQL. Compared to DQL, EDP proposes using DPM-Solver instead of DDPM to reduce the number of sampling steps. Then, EDP introduces an *action approximation* technique, where during policy optimization, one-step denoising is performed on the perturbed action a_t to approximate a_0. For the process using a data prediction model x_θ and a noise prediction model ϵ_θ separately, the following two equations can express the technique:

$$a_0 \approx x_\theta(a_t, t) \quad (34)$$

$$a_0 \approx \frac{a_t - \sigma_t \epsilon_\theta(a_t, t)}{\alpha_t}. \quad (35)$$

EDP reduces the sampling steps to 15 (even though DQL has only 5 sampling steps) and performs only one-step denoising during policy optimization, significantly speeding up the model training process and achieving performance close to that of DQL.

Implicit Diffusion Q-Learning. [16] Implicit Diffusion Q-Learning (IDQL) models the policy from the perspective of general *constrained policy search* (CPS), in which the optimal policy is described as a weighted behavior policy:

$$\pi^*(a|s) = \pi_\theta^b(a|s) w(a|s), \ s.t. \int_{\mathcal{A}} w(a|s) \mathrm{d}a = 1, \ \forall s, \quad (36)$$

where $\pi_\theta^b(a|s)$ represents the behavior policy learned by the diffusion model from the dataset, and $w(s, a)$ is a weight function. IDQL derives its weight function from the generalized implicit

12

Q-learning:

$$w(\boldsymbol{a}|\boldsymbol{s}) = \frac{|f'(Q_\phi(\boldsymbol{s},\boldsymbol{a}) - V^*(\boldsymbol{s}))|}{|Q_\phi(\boldsymbol{s},\boldsymbol{a}) - V^*(\boldsymbol{s})|}, \tag{37}$$

where f can be any convex function, $f' = \frac{\partial f}{\partial V(\boldsymbol{s})}$, and

$$V^*(\boldsymbol{s}) = \underset{V(\boldsymbol{s})}{\arg\min}\, \mathbb{E}_{\boldsymbol{a} \sim \pi_\theta^b(\boldsymbol{a}|\boldsymbol{s})}\left[f(Q_\phi(\boldsymbol{s},\boldsymbol{a}) - V(\boldsymbol{s}))\right]. \tag{38}$$

Therefore, the training of IDQL consists of two independent processes: training the diffusion model to clone the behavior policy and training the IQL-based weight function $w(\boldsymbol{a}|\boldsymbol{s})$. At each inference step, IDQL samples a set of candidate actions $\{\boldsymbol{a}_0\}$, computes the weights $\{w(\boldsymbol{s},\boldsymbol{a}_0)\}$, and then selects the action to be executed as a categorical from $\{w(\boldsymbol{s},\boldsymbol{a}_0)\}$.

DiffusionBC. [40] DiffusionBC constructs an observation-to-action diffusion model for imitating stochastic and multimodal human demonstrations. The basic version of DiffusionBC applies diffusion generation directly as a diffusion policy $\pi(\boldsymbol{a}_0|\boldsymbol{s},\boldsymbol{a}_t,t)$ with noisy action $\boldsymbol{a}_t \in \mathbb{R}^{|\boldsymbol{a}|}$, denoising timestep t and observation $\boldsymbol{s}$ (possibly with a history) input. To better select intra-distributional actions to mimic human behavior, DiffusionBC proposed the **Diffusion-X Sampling** trick, which encourages higher likelihood actions during sampling. For diffusion-X sampling, the sampling process first runs normal T denoising timesteps, and timesteps is fixed to $t = 1$, then extra denoising iterations continue to run for M timesteps toward higher-likelihood regions.

DiffusionPolicy. [4] Similar to DiffusionBC, Diffusion Policy also uses a diffusion model to directly approximate the conditional distribution $p(\boldsymbol{a}|\boldsymbol{s})$, but uses two key design choices: (1) **Closed-loop Action-chunking Prediction**: Diffusion Policy generates sequences of actions per prediction rather than single action to encourage temporal consistency and smoothness in long-term planning to better fit multimodal distributions. At time step t, the policy takes the latest T_s (the observation horizon) steps of observation data $\boldsymbol{s}_t$ as input and predicts H steps of actions, of which T_a (the action prediction horizon) steps of actions are executed on the robot without re-planning. (2) **Network Architecture Options**: Diffusion Policy adopts the traditional 1D-Unet [21] and DiT [41] to new CNN-based Unet and time-series diffusion transformer network architectures. CNN-based Diffusion Policy conditions the action generation process on observation $\boldsymbol{s}$ with Feature-wise Linear Modulation (FiLM) [42] and Transformer-based Diffusion Policy fuses state $\boldsymbol{s}$ and action $\boldsymbol{a}$ features via cross attention to jointly predict $\epsilon_\theta(o, a_k, k)$, where k is sinusoidal embedding for diffusion iteration. The Diffusion Policy has demonstrated excellent performance and high stability in multiple simulation environments and real-world tasks for imitation learning and is a widely used baseline for embodied AI.

G.3 Diffusion Data Synthesizers.

SynthER. [36] SynthER uses the diffusion model to generate one-step transitions $(\boldsymbol{s},\boldsymbol{a},r,d,\boldsymbol{s}')$. Trained on an offline dataset, SynthER then upsamples it to a larger dataset (in D4RL, SynthER upsamples each dataset to 5M transitions), which helps other offline RL algorithms to optimize the agent policy.

H Limitations, Challenges, and Future Directions

Limitations. Although the modular structure and pipeline design of `CleanDiffuser` greatly simplify the implementation difficulty for researchers deploying DMs, the inherent complexity of the principles and improvements of DMs still requires a considerable amount of time to deeply understand each type of module. We hope to alleviate this issue and better facilitate collaboration through comprehensive configuration files and documentation, as well as active maintenance and updates. Additionally, When dealing with certain specific issues, `CleanDiffuser` may require tailored adjustments and optimizations. For instance, the current version of `CleanDiffuser` does not directly support discrete or hybrid action space tasks, which may be mitigated through techniques such as action representation [28] or using categorical diffusion models [6].

13

Based on experimental analyses of `CleanDiffuser`, we have identified several promising areas for further research as follows:

Unleashing the potential of diffusion planners. Analogous to the classification of RL algorithms, as diffusion planners can imaginatively generate interactive trajectories, they should be categorized under model-based RL (MBRL). In MBRL, there are various ways to utilize learned dynamic models, including planning to search for the optimal action [13, 17], optimizing policies using rollout trajectories [12], and even combining these two approaches [15, 14]. Currently, diffusion planners are limited to the first paradigm, and due to their sensitivity to guidance and lack of safety constraints, they are prone to OOD plans [7], falling short in performance compared to other offline MBRL algorithms. Future research can explore new paradigms for diffusion planners, attempting diverse ways to utilize generated trajectories or integrating safety constraints to enhance the fidelity of generated trajectories, thereby unleashing the full potential of diffusion planners.

Exploring the reasons behind sampling degradation. In **??**, we discuss an anomaly known as *sampling degradation*, where the algorithm's performance decreases as the number of sampling steps increases. This anomaly has been identified in previous works [22, 3] and remains an open question. Theoretically, more sampling steps should result in a more accurate SDE/ODE solution, ultimately producing higher-fidelity samples. This naturally prompts a trade-off exploration between sampling steps and performance during implementation. However, in experiments, increasing sampling steps in certain tasks does not improve performance and can even lead to a decrease. Future research can systematically investigate this anomaly to provide optimal recommendations for selecting sampling steps.

Understanding the impact of SDE and ODE. In our experiments, we observe consistent differences in SDE solvers and ODE solvers on algorithm performance, tendency to sampling degradation, and sensitivity to guidance. While there is existing research on the impact of SDE and ODE in computer vision [39, 35], there is still a gap in research within the decision-making domain. Future research can fill this gap and explore the implications of SDE and ODE solvers in decision-making tasks.

Accelerating Diffusion Model Sampling. Due to the denoising process involved in iterative sampling, DMs face the issue of slow sampling speeds when used for decision-making. This poses significant challenges in scenarios such as real-time robot control or game AI. DiffuserLite [7] is a diffusion planner method that addresses this issue by modeling the diffusion process through a plan refinement process for coarse-to-fine-grained trajectory generation and further accelerates the sampling speed using rectified flow. Further speeding up the sampling speed of various roles of DMs remains a promising research direction.

I Potential Social Impact

`CleanDiffuser` fills a critical gap in the current landscape by providing a unified and modularized framework that empowers researchers and practitioners to explore new frontiers. This will accelerate the development and deployment of diffusion-based decision-making applications, such as various robotics research and products. However, `CleanDiffuser` may also be used in military weapon development.

J License

Our codebase is released under Apache License 2.0.

14

References

[1] Anurag Ajay, Yilun Du, Abhi Gupta, Joshua B. Tenenbaum, Tommi S. Jaakkola, and Pulkit Agrawal. Is conditional generative modeling all you need for decision making? In The Eleventh International Conference on Learning Representations, ICLR, 2023.

[2] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. arXiv preprint arXiv:1606.01540, 2016.

[3] Huayu Chen, Cheng Lu, Chengyang Ying, Hang Su, and Jun Zhu. Offline reinforcement learning via high-fidelity generative behavior modeling. In The Eleventh International Conference on Learning Representations, ICLR, 2023.

[4] Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In Proceedings of Robotics: Science and Systems, RSS, 2023.

[5] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In Advances in Neural Information Processing Systems, NIPS, 2021.

[6] Sander Dieleman, Laurent Sartran, Arman Roshannai, Nikolay Savinov, Yaroslav Ganin, Pierre H Richemond, Arnaud Doucet, Robin Strudel, Chris Dyer, Conor Durkan, et al. Continuous diffusion for categorical data. arXiv preprint arXiv:2211.15089, 2022.

[7] Zibin Dong, Jianye Hao, Yifu Yuan, Fei Ni, Yitian Wang, Pengyi Li, and Yan Zheng. Diffuserlite: Towards real-time diffusion planning. arXiv preprint arXiv:2401.15443, 2024.

[8] Pete Florence, Corey Lynch, Andy Zeng, Oscar A Ramirez, Ayzaan Wahid, Laura Downs, Adrian Wong, Johnny Lee, Igor Mordatch, and Jonathan Tompson. Implicit behavioral cloning. In Conference on Robot Learning, CoRL, 2022.

[9] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4rl: Datasets for deep data-driven reinforcement learning. arXiv preprint arXiv:2004.07219, 2020.

[10] Scott Fujimoto and Shixiang Gu. A minimalist approach to offline reinforcement learning. In Advances in Neural Information Processing Systems, NIPS, 2021.

[11] Abhishek Gupta, Vikash Kumar, Corey Lynch, Sergey Levine, and Karol Hausman. Relay policy learning: Solving long-horizon tasks via imitation and reinforcement learning. In Proceedings of the Conference on Robot Learning, CoRL, 2020.

[12] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. In International Conference on Learning Representations, ICLR, 2020.

[13] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In Proceedings of the 36th International Conference on Machine Learning, ICML, 2019.

[14] Nicklas Hansen, Hao Su, and Xiaolong Wang. TD-MPC2: Scalable, robust world models for continuous control. In The Twelfth International Conference on Learning Representations, ICLR, 2024.

[15] Nicklas A Hansen, Hao Su, and Xiaolong Wang. Temporal difference learning for model predictive control. In Proceedings of the 39th International Conference on Machine Learning, ICML, 2022.

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

15

[17] Xiaotian Hao, Jianye Hao, Chenjun Xiao, Kai Li, Dong Li, and Yan Zheng. Multiagent gumbel muzero: Efficient planning in combinatorial action spaces. Proceedings of the AAAI Conference on Artificial Intelligence, AAAI, 2024.

[18] Longxiang He, Li Shen, Linrui Zhang, Junbo Tan, and Xueqian Wang. Diffcps: Diffusion model based constrained policy search for offline reinforcement learning. arXiv preprint arXiv:2310.05333, 2024.

[19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In Advances in Neural Information Processing Systems, NIPS, 2020.

[20] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications, 2021.

[21] Michael Janner, Yilun Du, Joshua Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. In Proceedings of the 39th International Conference on Machine Learning, ICML, 2022.

[22] Bingyi Kang, Xiao Ma, Chao Du, Tianyu Pang, and Shuicheng Yan. Efficient diffusion policies for offline reinforcement learning. Advances in Neural Information Processing Systems, NIPS, 36, 2024.

[23] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, Advances in Neural Information Processing Systems, NIPS, 2022.

[24] Diederik P Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. In Advances in Neural Information Processing Systems, NIPS, 2021.

[25] P.E. Kloeden and E. Platen. Numerical Solution of Stochastic Differential Equations. Stochastic Modelling and Applied Probability. Springer Berlin Heidelberg, 2011.

[26] Ilya Kostrikov, Ashvin Nair, and Sergey Levine. Offline reinforcement learning with implicit q-learning. In International Conference on Learning Representations, ICLR, 2022.

[27] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. In Advances in Neural Information Processing Systems, NIPS, 2020.

[28] Boyan Li, Hongyao Tang, Yan Zheng, Jianye Hao, Pengyi Li, Zhen Wang, Zhaopeng Meng, and Li Wang. Hyar: Addressing discrete-continuous action reinforcement learning via hybrid action representation. arXiv preprint arXiv:2109.05490, 2021.

[29] Wenhao Li, Xiangfeng Wang, Bo Jin, and Hongyuan Zha. Hierarchical diffusion for offline decision making. In Proceedings of the 40th International Conference on Machine Learning, ICML, 2023.

[30] Eric Liang, Richard Liaw, Robert Nishihara, Philipp Moritz, Roy Fox, Ken Goldberg, Joseph Gonzalez, Michael Jordan, and Ion Stoica. RLlib: Abstractions for distributed reinforcement learning. In Proceedings of the 35th International Conference on Machine Learning, ICML, 2018.

[31] Zhixuan Liang, Yao Mu, Mingyu Ding, Fei Ni, Masayoshi Tomizuka, and Ping Luo. Adaptdiffuser: Diffusion models as adaptive self-evolving planners. In International Conference on Machine Learning, ICML, 2023.

[32] Haohe Liu, Zehua Chen, Yi Yuan, Xinhao Mei, Xubo Liu, Danilo Mandic, Wenwu Wang, and Mark D Plumbley. AudioLDM: Text-to-audio generation with latent diffusion models. In Proceedings of the 40th International Conference on Machine Learning, ICML, 2023.

16

[33] Xingchao Liu, Chengyue Gong, and qiang liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In The Eleventh International Conference on Learning Representations, ICLR, 2023.

[34] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. DPM-solver: A fast ODE solver for diffusion probabilistic model sampling in around 10 steps. In Advances in Neural Information Processing Systems, NIPS, 2022.

[35] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. arXiv preprint arXiv:2211.01095, 2023.

[36] Cong Lu, Philip Ball, Yee Whye Teh, and Jack Parker-Holder. Synthetic experience replay. Advances in Neural Information Processing Systems, NIPS, 36, 2024.

[37] Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. In 5th Annual Conference on Robot Learning, CoRL, 2021.

[38] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In Proceedings of the 38th International Conference on Machine Learning, ICML, 2021.

[39] Shen Nie, Hanzhong Allan Guo, Cheng Lu, Yuhao Zhou, Chenyu Zheng, and Chongxuan Li. The blessing of randomness: SDE beats ODE in general diffusion-based image editing. In The Twelfth International Conference on Learning Representations, ICLR, 2024.

[40] Tim Pearce, Tabish Rashid, Anssi Kanervisto, Dave Bignell, Mingfei Sun, Raluca Georgescu, Sergio Valcarcel Macua, Shan Zheng Tan, Ida Momennejad, Katja Hofmann, and Sam Devlin. Imitating human behaviour with diffusion models. In The Eleventh International Conference on Learning Representations, ICLR, 2023.

[41] William Peebles and Saining Xie. Scalable diffusion models with transformers. In Proceedings of the IEEE/CVF International Conference on Computer Vision, ICCV, 2023.

[42] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In Proceedings of the AAAI conference on artificial intelligence, AAAI, 2018.

[43] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-baselines3: Reliable reinforcement learning implementations. Journal of Machine Learning Research, 2021.

[44] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR, 2022.

[45] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR, 2023.

[46] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In International Conference on Learning Representations, ICLR, 2021.

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

17

[48] Denis Tarasov, Alexander Nikulin, Dmitry Akimov, Vladislav Kurenkov, and Sergey Kolesnikov. CORL: Research-oriented deep offline reinforcement learning library. In 3rd Offline RL Workshop: Offline RL as a "Launchpad", 2022.

[49] Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, Dhruv Nair, Sayak Paul, William Berman, Yiyi Xu, Steven Liu, and Thomas Wolf. Diffusers: State-of-the-art diffusion models. `https://github.com/huggingface/diffusers`, 2022.

[50] Zhendong Wang, Jonathan J Hunt, and Mingyuan Zhou. Diffusion policies as an expressive policy class for offline reinforcement learning. In The Eleventh International Conference on Learning Representations, ICLR, 2023.

[51] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In Proceedings of the IEEE/CVF International Conference on Computer Vision, ICCV, 2023.

[52] Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware. In Proceedings of Robotics: Science and Systems, RSS, 2023.

18