

Crossway Diffusion: Improving Diffusion-based Visuomotor Policy via Self-supervised Learning

Xiang Li, Varun Belagali, Jinghuan Shang and Michael S. Ryoo

Abstract—Diffusion models have been adopted for behavioral cloning in a sequence modeling fashion, benefiting from their exceptional capabilities in modeling complex data distributions. The standard diffusion-based policy iteratively denoises action sequences from random noise conditioned on the input states and the model is typically trained with a singular diffusion loss. This paper explores the potential enhancements in such models when the denoising process is informed by a better visual representation. We study the scenario where the model is jointly optimized using the standard diffusion loss alongside an auxiliary objective based on self-supervised learning. After experimenting with various objectives, we introduce Crossway Diffusion, a simple yet effective way to enhance diffusion-based visuomotor policy learning via a state decoder and an auxiliary reconstruction objective. During training, the state decoder reconstructs raw image pixels and other states from the intermediate representations of the model. Experiments demonstrate the effectiveness of our method in various simulated and real-world tasks, confirming its consistent advantages over the standard diffusion-based policy and other baselines.

I. INTRODUCTION

Behavioral Cloning (BC) [1] is a supervised learning formulation for robot action policy learning. Given expert demonstration data consisting of a sequence of state-action pairs, a model is trained to predict the correct action vector given input states (e.g., images). This framework has shown to be very effective particularly when a sufficient amount of training data is provided [2], [3].

Recently, sequence modeling approaches [4], [5], [6] have been often used for behavioral cloning, because of their ability to model multiple steps of information. In such a formulation, the objective is to model the probability distribution of the multi-step state-action trajectory. This allows BC to consider beyond a single-step regression, better-taking advantage of history. Given the success in modeling human language [7] and images [8], Transformers [9] have been popularly adopted for sequence modeling-based policies [4], [5], [6], [2], [10].

Diffusion models [11], [12], [13], [14] have exceptional capabilities in modeling multimodal data distribution and generating new samples from the distribution, which make them suitable for imitating behaviors by generating trajectories. Recent works [15], [16], [17] have successfully applied diffusion models for sequential modeling using low dimensional states. For visuomotor control tasks, [18] demonstrated

The authors are with the Department of Computer Science, Stony Brook University, Stony Brook, New York, 11790. {xiangli8, vbelagali, jishang, mryoo}@cs.stonybrook.edu

The appendix, code, pretrained checkpoints, and datasets are available at https://github.com/LostXine/crossway_diffusion.

promising performance using multimodal states including visual observations as the conditions of the diffusion model.

In this work, we explore the potential enhancements in diffusion-based visuomotor policy when conditioning on a better visual representation. Particularly, we study the scenario where the model is jointly optimized using the standard diffusion loss alongside an auxiliary objective based on self-supervised learning (SSL). After extensively experimenting with various objectives, we introduce Crossway Diffusion, a simple yet effective way to enhance diffusion-based visuomotor policy learning via a state decoder and an auxiliary reconstruction objective. During training, the state decoder reconstructs raw image pixels and other states from the intermediate representations of the model. By doing so, we explicitly regularize the intermediate representation to capture the information of the input states, as now the model needs to reconstruct the input states from it.

Through our experiments over multiple challenging tasks from different benchmarks, we verify the consistent advantage of our method in comparison to the baseline Diffusion Policy [18]. Especially, our method achieves a 15.7% improvement in success rate over the baseline on ‘Transport, mh’ dataset from Robomimic [19] and one variant of our method even achieves 17.1%, emphasizing the effectiveness of the SSL objective and the specific state decoder design. Our contribution can be summarized as follows:

- We delve into the challenge of improving diffusion-based visuomotor policy learning via jointly optimizing the diffusion loss and an auxiliary SSL objective
- We propose Crossway Diffusion, which consistently enhances diffusion-based visuomotor policy learning via a state decoder and an auxiliary reconstruction objective
- We confirm the effectiveness of the proposed method on multiple challenging visual BC tasks from different benchmarks, including two real-world robot manipulation datasets we collected.
- We conduct detailed ablations on multiple design choices, verifying the advance and robustness of the proposed design.

II. PRELIMINARIES

A. Behavioral Cloning

We consider simple behavioral cloning (BC) setting over a Markov Decision Process (MDP), described by the tuple $(\mathcal{S}, \mathcal{A}, P)$, where $s \in \mathcal{S}$ represents the state, $a \in \mathcal{A}$ is the action, and P are the transition dynamics given by $P(s'|s, a)$. A trajectory consists of a sequence of state-action

pairs $\{s_0, a_0, s_1, a_1, \dots, s_T, a_T\}$, where T is the length of sequence (task horizon). Our goal is to train a robot policy π that best recovers an unknown policy π^* using a demonstration dataset $\mathcal{D} = \{(s_i, a_i)\}$ collected by π^* . Specifically, the robot policy π operates on a trajectory basis: $\pi(A_t|S_t)$, where $S_t = \{s_{t-T_s+1}, s_{t-T_s+2}, \dots, s_t\}$ is the given short history state sequence, $A_t = \{a_t, a_{t+1}, \dots, a_{t+T_a-1}\}$ is the predicted future actions to take. T_s and T_a represent the lengths of these two sequences respectively.

B. Diffusion Models

Diffusion models [11], [12], [13] are generative models that iteratively generate samples that match the data distribution. Recent works [20], [21], [22], [23] have demonstrated its great ability in image generation and other data generation tasks. The diffusion process is of the original data being destroyed by a sequence of noise $q(x^k|x^{k-1})$ known as the forward process, where k is the current diffusion step and there are K steps in total. The diffusion model uses the reverse process, the backward process, $p_\theta(x^{k-1}|x^k)$ to denoise the corrupted data. By iteratively denoising, a diffusion model generates synthesized data \hat{x} that approximates the original data distribution $q(x_0)$, starting from a random prior $p(x^K)$:

$$p_\theta(x^0) = \int p_\theta(x^{0:K}) dx^{1:K} \quad (1)$$

$$= \int p(x^K) \prod_{k=1}^K p_\theta(x^{k-1}|x^k) dx^{1:K} \quad (2)$$

Typically the random prior $p(x^K)$ is a standard Gaussian distribution, and the denoising process is parameterized by the following Gaussian:

$$p_\theta(x^{k-1}|x^k) = \mathcal{N}(x^{k-1}|\mu_\theta(x^k, k), \Sigma^k), \quad (3)$$

where the parameter $\mu_\theta(x^k, k)$ is estimated from a neural network parameterized by θ .

C. Diffusion Models for Policy Learning

Diffusion models have been applied for data augmentation [24] and sample synthesis [25], [26], [27] in robot policy learning. As recent works [4], [5], [28] formulate the robot policy learning as an action sequence generation problem, diffusion models have also been successfully adapted as sequence generation models [17], [24], [16], [15]. Diffuser [15] concatenates the low-dimensional states and the actions and models multiple state-action pairs as a matrix (image) generation problem. However, such a setting is not feasible for visuomotor policy learning due to the high dimensionality of the visual observations.

Diffusion Policy [18] tackles the challenge by generating only action sequences, while conditioned on visual and other states. Specifically, given a sequence of T_s states $S_t = \{s_{t-T_s+1}, s_{t-T_s+2}, \dots, s_t\}$ where each s contains both visual states and low-dimensional states, the diffusion model generates a sequence of T_a actions $A_t = \{a_t, a_{t+1}, \dots, a_{t+T_a-1}\}$ conditioned on the state sequence. When the agent finishes




Fig. 1: Trajectory generation formulation of Diffusion Policy. This figure shows a case where $T_s = 2$ and $T_a = 4$.

executing A_t , Diffusion Policy generates the consequent action sequence A_{t+T_a} given S_{t+T_a} , which formulates a closed-loop control (see Fig. 1).

The model of Diffusion Policy is composed of a state encoder E_S , an action encoder E_A , and an action decoder D_A . The action encoder and decoder make up the diffusion model for generating action sequences by running the denoising process iteratively. The state encoder provides conditioning from the states, which modulates the generation process. This paper focuses on the convolutional version of the Diffusion Policy due to its superior performance.

Given a state sequence with both visual and low-dimensional states $S_t = \{S_{t,img}, S_{t,low-dim}\}$, the state encoder extracts visual embeddings from images $h_{t,img} = E_S(S_{t,img})$. The visual embeddings $h_{t,img}$ are then concatenated with other low dimensional states $S_{t,low-dim}$ to form the observation condition $h_t = h_{t,img} \oplus S_{t,low-dim}$.

The action encoder E_A takes the noisy action sequence A_t^k at diffusion step k and the observation condition h_t and produces the representation X_t^k from the deepest layer and other tensors for skip connection $X_{t,skip}^k$ from the shallower layers: $X_t^k, X_{t,skip}^k = E_A(A_t^k, h_t, k)$, where t is the timestamp of a state or action trajectory. $X_t^k \in \mathbb{R}^{T \times C}$ where T is the representation length along the time axis and C is the number of channels.

The action decoder takes both $X_t^k, X_{t,skip}^k$ and condition h_t to estimate the noise ϵ which is applied to A_t^k in the forward diffusion process. The condition h_t is applied between two convolutional layers in the residual block, using Feature-wise Linear Modulation (FiLM) [29]. Then a (slightly) denoised action sequence A_t^{k-1} is derived from the estimated noise ϵ_θ and A_t^k using Eq.4.

$$\begin{aligned} \epsilon_\theta &= D_A(X_t^k, X_{t,skip}^k, h_t) \\ A_t^{k-1} &= \frac{1}{\sqrt{\alpha_k}} \left(A_t^k - \frac{1 - \alpha_k}{\sqrt{1 - \alpha_k}} \epsilon_\theta \right) + \sigma_k z \end{aligned} \quad (4)$$

where z is randomly sampled from a normal distribution with the same dimension as A_t^k . $\alpha_k, \hat{\alpha}_k$, and σ_k are parameters regarding the diffusion process used in DDPM [12], except that we use k as the diffusion step instead of t .

During inference, the denoising process mentioned above is repeated for K times iteratively, generating a noiseless action sequence A_t^0 in the end. The whole model is optimized using the same Mean Squared Error (MSE) as DDPM [12] to predict the noise ϵ applied to A_t^0 for constructing A_t^k .

$$\mathcal{L}_{DDPM} = \text{MSE}(D_S(E_A(A_t^k, h_t, k), h_t), \epsilon) \quad (5)$$

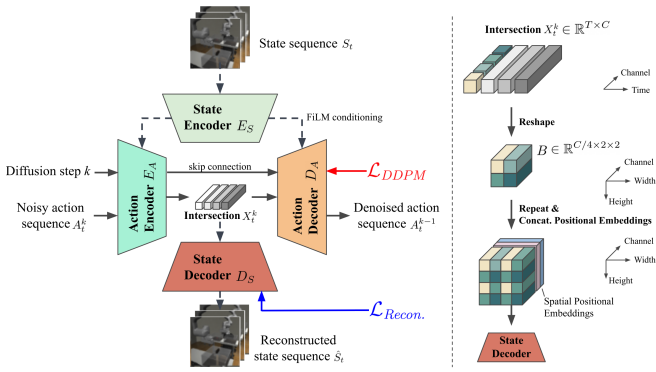


Fig. 2: Left: Architecture of Crossway Diffusion. We introduce a state decoder to the existing Diffusion Policy [18] as well as an auxiliary reconstruction objective \mathcal{L}_{Recon} . The state decoder takes a transformed intermediate representation ‘intersection’ to reconstruct the input states. Right: Transformation applied to ‘intersection’ for the visual state decoder.

III. METHOD

Crossway Diffusion extends existing Diffusion Policy [18] by introducing (1) a state decoder and (2) an auxiliary objective, both for reconstructing the input states during training. The overall architecture of Crossway Diffusion is presented as Fig. 2. Specifically, the state decoder takes the intermediate representation of the diffusion process X_t^k to reconstruct the input states, and the reconstruction objective is jointly optimized with the diffusion loss \mathcal{L}_{DDPM} .

A. State Decoder

The state decoder D_S is newly introduced to reconstruct the input states from the intermediate representation along with the existing action sequence denoising pipeline. The new procedure is summarized as $\hat{S}_t = D_S(g(X_t^k))$, where \hat{S}_t is the reconstructed state and $g(\cdot)$ is the *intersection transformation*, discussed in the following section.

The representation X_t^k is dubbed as *intersection* since both the flow of action denoising and the flow of state reconstruction pass through this tensor and then head to their corresponding destinations. For the same reason, we name our method *Crossway Diffusion*, whose key feature is to have two flows of information intersected in the diffusion model during the training phase. By doing so, we explicitly regularize X_t^k to capture the information from both flows, which benefits representation learning. Our design includes two distinct types of state decoders: one for visual states and another for low-dimensional states. For each source of the states (multiple cameras, joint angles, and so on), we assign a dedicated decoder for the best reconstruction results.

To reconstruct the visual states, the visual state decoders are made of a sequence of 2D residual convolutional blocks, transposed convolutional layers for upsampling (ConvTranspose), and vanilla convolutional layers (See Fig. 3). The numbers in the blocks indicate the number of output channels except for ConvTranspose. ConvTranspose doubles the spatial resolution of the input tensor while keeping the number




Fig. 3: Architecture of visual state decoder. Numbers in the blocks indicate the number of output channels except that ConvTranspose doubles the spatial resolution while keeping the number of channels unchanged.

of channels unchanged. Inspired by [30] and [31], we adapt the positional embedding from [32] to present the pixel coordinates. The positional embedding is concatenated to the output of ConvTranspose along the channel axis.

Other low-dimensional states are regressed by three-layer MLPs. The widths of the hidden layers are in the ratios of 4:2:1 compared to the width of the low-dimensional states.

Note that the reconstruction with the state decoder D_S is only used during the training, serving as an ‘interpreter’ to generate additional supervisory signals to train better intermediate representations. The state decoder will be discarded after the training.

B. Intersection Transformation

Intersection transformation $g(\cdot)$ adapts the intermediate representation X_t^k for the state decoder to pursue the reconstruction task. Assume X_t^k is split along the time axis as a list of vectors $X_t^k := \{X_{t,i}^k | i = 1, 2, \dots, T\}$. By default, only the first vector in the list $X_{t,1}^k$ is sent to the state decoder for state reconstructions.

The transformation for the low-dimensional state decoder $g_{low-dim}(\cdot)$ is an identity function which means the state decoder directly takes $X_{t,1}^k$ as the input.

In contrast, the visual state decoder employs another transformation $g_{img}(\cdot)$ due to the significant dimensional disparity between $X_{t,1}^k$ and the visual states. As shown in the right of Fig. 2, to bridge the dimensionality gap, the C elements of $X_{t,1}^k$ are equally split into 4 folds and then reshaped as a $C/4 \times 2 \times 2$ block B . The block B is then repeated multiple times in two spatial dimensions so that it will have a spatial resolution of a quarter of the desired reconstructed image along each spatial dimension. Finally, we concatenate the same positional embeddings introduced in the previous section along the channel axis before sending the transformed intersection to the visual state decoder.

We further explore the effect of multiple designs of $g(\cdot)$ in Section IV-C.

C. Crossway Diffusion Loss

Given the state decoder and the state reconstruction task, we introduce a reconstruction loss \mathcal{L}_{Recon} between the reconstructed states and the original input states, which provides an auxiliary training signal to D_S , E_S , and E_A . \mathcal{L}_{Recon} is simply implemented as a Mean Squared Error (MSE).

In addition to \mathcal{L}_{DDPM} used in Diffusion Policy [18] (Eq. 5), which supervises D_A , E_A and E_S , \mathcal{L}_{Recon} is jointly optimized with \mathcal{L}_{DDPM} by a simple weighted addition. That is, all network modules E_A , D_A , E_S , and E_S are trained jointly.

The total loss for Crossway Diffusion is denoted as Eq.6 and we find $\alpha = 0.1$ is a generally good setting without extensive hyperparameter search.

$$\begin{aligned} \mathcal{L}_{Recon.} &= \text{MSE}(S_t, \hat{S}_t) \\ \mathcal{L}_{Crossway} &= \mathcal{L}_{DDPM} + \alpha \mathcal{L}_{Recon.} \end{aligned} \quad (6)$$

IV. EXPERIMENT

We first evaluate Crossway Diffusion on multiple challenging simulated tasks and two real-world tasks. Then, we explore variants of Crossway Diffusion regarding state decoder design and the auxiliary objectives. Through an extensive amount of experiments, we confirm that Crossway Diffusion consistently leads to better performance than vanilla Diffusion Policy [18] and other baselines [33] over all the tested tasks.

A. Task and Dataset

We follow [18] and choose five tasks Can, Lift, Square, Transport, and Tool Hang from Robomimic [19] and Push-T from Implicit Behaviour Cloning (IBC) [33]. In addition, we build a real-world robot arm manipulation environment and collect our own data for two tasks.

In the ‘Can’ task, the robot needs to lift a soda can from one box and put it into another box. In the ‘Lift’ task, the robot needs to lift a cube above a certain height. In the ‘Square’ task, the robot needs to fit the square nut onto the square peg. The ‘Transport’ task entails the collaborative effort of two robot arms to transfer a hammer from a closed container on one table to a bin on another table. One arm is responsible for retrieving and passing the hammer, while the other arm cleans the bin and receives the passed hammer. In ‘Tool Hang’, the robot needs to insert the hook into the base to assemble a frame and then hang a wrench on the hook. Additionally, the ‘Push-T’ task involves pushing a T-shaped block (gray) onto a target location (green) in a 2D space. Both two types of datasets are evaluated if available: ‘ph’ - proficient-human demonstration and ‘mh’ - multi-human demonstrations, originally defined by Robomimic [19]. ‘mh’ is designed to have diverse proficiency on the task.

In the real-world environment, we set up a robot arm with a gripper and two RGB webcams. The first camera is stationary and offers a third-person perspective of the operating space, while the second camera is mounted on the gripper, providing a first-person view for grasping. The action space encompasses both the 3D position of the robot arm and a binary signal for the gripper’s opening and closing. For task ‘Duck Lift’, the objective is to lift a rubber duck above a certain height. For task ‘Duck Collect’, the robot needs to collect four ducks and sort them into two separate containers based on the colors of the ducks.

A visual reference for all the tasks is provided as Fig. 4 and an information summary is presented in Tab. I.

B. Main Results

1) *Evaluation Metrics:* Consistent with prior studies, for the Push-T task, we measure the extent of target location coverage achieved by the T block, which is the ratio of the




Fig. 4: Visual reference for all tasks

TABLE I: Dataset summary. ph: the number of proficient-human demonstrations; mh: the number of multi-human demonstrations; R?: whether the dataset is a real-world dataset or not; Rob.: the number of robots; Obj.: the number of objects; Cam.: the number of cameras; Act-D: action dimension; Steps: max number of rollout steps.

Task	ph	mh	R?	Rob.	Obj.	Cam.	Act-D	Steps
Can	200	300	N	1	1	2	7	400
Lift	200	300	N	1	1	2	7	400
Square	200	300	N	1	1	2	7	400
Transport	200	300	N	2	3	4	14	700
Tool Hang	200	-	N	1	2	2	7	700
Push-T	200	-	N	1	1	1	2	300
Duck Lift	100	-	Y	1	1	2	4	50
Duck Collect	100	-	Y	1	1	2	4	200

covered area to the total area. We adopt the success rate of other tasks as a performance metric. All the models are trained for 500 epochs and evaluated at the end of training. Note that we train and evaluate all methods from scratch with our setting for a fair comparison, as the numbers from [18] are not reliable because of the bug in their evaluation code¹. For all diffusion-based methods, we benchmark the exponential moving average (EMA) version of the model for better stability, as suggested by [12]. More training details are available in Section VI-A.

2) *Simulated Experiments:* For all the simulated tasks, in Tab. II we report the average performance over 1000 randomly initialized episodes \times 3 models trained with random seeds, as well as the standard deviation among random seeds. That is, each score before \pm in Tab. II is an average of 3000 episodes, and each number after \pm is the standard deviation of the score of 3 random seeds.

From the comparison, the proposed Crossway Diffusion consistently outperforms the baseline Diffusion Policy [18] in all datasets, as well as other baselines. We observe an improvement of 15.7% over the success rate in ‘Transport, mh’, emphasizing the effectiveness of our method when the demonstration data is varied in proficiency. Please refer to Section VI-G and VI-D for our example episodes and visualization of the action generation process.

3) *Real-world Experiments:* For all tested methods, we run ‘Duck Lift’ for 20 episodes and ‘Duck Collect’ for 10

¹Please check this link for more details https://github.com/real-stanford/diffusion_policy/issues/6.

TABLE II: Scores on simulated datasets. We report the average of 3000 episodes and the standard deviation of 3 seeds

Method	Can, ph	Can, mh	Lift, ph	Lift, mh	Square, ph	Square, mh	Transport, ph	Transport, mh	Tool Hang, ph	Push-T
LSTM-GMM	0.714 ± 0.247	0.887 ± 0.033	0.978 ± 0.017	0.992 ± 0.001	0.643 ± 0.023	0.491 ± 0.057	0.656 ± 0.049	0.254 ± 0.017	0.460 ± 0.060	0.567 ± 0.013
IBC [33]	0.008 ± 0.006	0.001 ± 0.001	0.709 ± 0.008	0.222 ± 0.112	0.002 ± 0.001	0.000 ± 0.001	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000	0.687 ± 0.031
Diffusion Policy CNN [18]	0.992 ± 0.002	0.958 ± 0.003	1.000 ± 0.000	0.998 ± 0.001	0.935 ± 0.006	0.858 ± 0.007	0.859 ± 0.015	0.643 ± 0.004	0.772 ± 0.012	0.819 ± 0.002
Crossway Diffusion (Ours)	0.994 ± 0.002	0.965 ± 0.003	1.000 ± 0.000	0.998 ± 0.000	0.935 ± 0.005	0.879 ± 0.010	0.864 ± 0.016	0.800 ± 0.020	0.792 ± 0.014	0.843 ± 0.020

TABLE III: Success rate of real-world tasks

	Duck Lift	Duck Collect
Diffusion Policy CNN [18]	0.80	0.70
Crossway Diffusion (Ours)	0.95	0.80




Fig. 5: Duck Lift task under different obstructions. Two rows show the images captured by two cameras respectively. The red arrow in (e) is used to indicate the position of the duck.

episodes and measure the success rate. Note that in task ‘Duck Collect’, one episode is considered a success only if the robot successfully collects and sorts all four ducks. For each episode, we place the ducks at random initial positions but keep the positions consistent across tested methods. The results are reported in Tab. III, highlighting the advantages of our method. We also show that our method is robust to distractions like unseen objects, occlusions in one camera, and other distractions (Fig. 5(a-d)). Please refer to Section VI-E and the video for more information.

Additionally, we provide qualitative results on image reconstruction in Section VI-F, where the quality of reconstructed visual states is surprisingly high.

C. Ablations

The ablation study mainly focuses on two critical designs in our method: the state decoder and the auxiliary objective.

1) *On State Decoder*: Experiments in this section are designed to answer (1) which representation for state reconstruction benefits policy learning most, as well as (2) what is the best architecture for the visual state decoder.

First, two more designs (Design B and C) of the intersection transformation presented in Fig. 6 are studied over five simulated datasets, as well as one variant (Design D) that utilizes h_t instead of intersection X_t^k for reconstruction.

Design A is the default Crossway Diffusion introduced in Section III. In Design B, the first $C/2$ channels for all vectors in intersection X_t^k are selected for reconstruction,




Fig. 6: Selected parts for state reconstruction in three designs. White and gray parts are not used for reconstruction.

TABLE IV: Ablations on intersection transformation.

	Square, mh	Transport, ph	Transport, mh	Tool Hang, ph	Push-T
A (default)	0.879 ± 0.010	0.864 ± 0.016	0.800 ± 0.020	0.792 ± 0.014	0.843 ± 0.020
B	0.881 ± 0.017	0.882 ± 0.010	0.784 ± 0.025	0.777 ± 0.010	0.835 ± 0.012
C	0.868 ± 0.006	0.906 ± 0.012	0.814 ± 0.028	0.783 ± 0.005	0.831 ± 0.003
D	0.873 ± 0.012	0.892 ± 0.002	0.764 ± 0.013	0.790 ± 0.007	0.819 ± 0.015
Diff. [18]	0.858 ± 0.007	0.859 ± 0.015	0.643 ± 0.004	0.772 ± 0.012	0.819 ± 0.002

while the rest is left dedicated to the denoising process. In contrast, Design C takes advantage of all vectors in X_t^k for the reconstruction. Additionally, for both Design B and C, the selected vectors are independently projected by a linear layer to match the target number of channels $C/4$. The latter operations like reshape and repeat in Fig. 2 are kept the same. Finally, in Design D, the state decoder directly takes the output of the state encoder h_t instead of X_t^k , which disentangles the action denoising flow and the state reconstruction flow. A linear layer is also applied to project h_t to the target dimension $C/4$ in the channel axis.

The results presented in Tab. IV show that Design A, B, and C consistently outperform the baseline Diffusion Policy CNN [18] (Diff. in table). However, Design D achieves similar results to the baseline on Push-T and shows improvements on other datasets. Such observations validate the importance of not only the *reconstruction task* but also the design which *explicitly forces two flows to intersect with each other*. Though these designs show diverse advantages over different tasks, we choose Design A as the default due to its computational simplicity.

Then, the variants with different visual state decoder architectures are studied. The default visual state decoder contains two stages where each stage has two residual blocks with a transposed convolution for spatial upsampling (see Fig. 3). *Shallower Dec.* takes the second stage out and only keeps the first stage. For *ViT Dec.*, the visual state decoder is a two-layer Vision Transformer (ViT) [8] whose token dimension is $C/4$ and the positional embeddings are directly added to the tokens before the ViT instead of concatenation.

The results on Push-T are presented in Tab. V (more available in Appendix VI-B), showing that all variants of our methods perform better than the baseline Diffusion Policy CNN [18] (Diff. in table). Such observation validates the effectiveness of the reconstruction task and our specific visual state decoder design.

2) *On auxiliary objective*: By default, both visual states and low-dimensional states are reconstructed in Crossway

TABLE V: Ablation results tested on Push-T

Default	Shallower Dec.	ViT Dec.	Visual-only	Diff. [18]
0.843 ± 0.020	0.822 ± 0.014	0.824 ± 0.008	0.828 ± 0.012	0.819 ± 0.002

TABLE VI: Ablations on future prediction, tested on Push-T

$N = 0$ (default)	$N = 2$	$N = 4$	$N = 6$	$N = 8$
0.843 ± 0.020	0.818 ± 0.006	0.827 ± 0.014	0.817 ± 0.003	0.803 ± 0.013

TABLE VII: Performance of Crossway-CURL, which adopts a contrastive loss as the auxiliary objective

	Lift, mh	Lift, ph	Square, mh	Square, ph	Push-T
Crossway-CURL	0.802 ± 0.024	0.678 ± 0.188	0.053 ± 0.025	0.035 ± 0.007	0.518 ± 0.160
Default	0.998 ± 0.000	1.000 ± 0.000	0.879 ± 0.010	0.935 ± 0.005	0.843 ± 0.020
Diff. [18]	0.998 ± 0.001	1.000 ± 0.000	0.858 ± 0.007	0.935 ± 0.006	0.819 ± 0.002

Diffusion. We first benchmark a simple variant called *Visual-only*, which reconstructs only the visual states. Results in Tab. V verify the benefit of predicting all the input states.

Then we test the variants where the state decoder predicts the state that is N steps ahead of the current state, instead of predicting (reconstructing) the current state as in the default setting ($N = 0$). From Tab. VI we empirically learn that reconstructing the current state is the most beneficial method while predicting future states may compromise the performance, even worse than the baseline. We further extend the study to multiple challenging datasets and obtain a similar observation included in Appendix VI-B.

Furthermore, we verify the effectiveness of the proposed reconstruction task in comparison to another SSL objective using contrastive learning inspired by CURL [34]. In particular, we independently perform random crop on all images of an observation sequence S_t twice to get two augmented sequences $S_{t,a1}$ and $S_{t,a2}$. The model takes $S_{t,a1}$ to produce intersection $X_{t,a1}^k$ while $S_{t,a2}$ is processed by the exponential moving average (EMA) version of the model to get $X_{t,a2,ema}^k$. The intuition is that due to the semantic similarity between $S_{t,a1}$ and $S_{t,a2}$, the intersection $X_{t,a1}$ and $X_{t,a2,ema}$ should also be similar in a latent space.

We follow CURL [34] and maintain a learnable matrix W . For each batch of b samples, we calculate the similarity matrix M_{sim} between all samples in the same batch using matrix multiplication (first line in Eq.7, where $sg(\cdot)$ means stop gradients). Then the contrastive loss \mathcal{L}_{CURL} is formulated as Eq. 7, where b is the batch size and $\alpha = 0.1$. Finally \mathcal{L}_{CURL} is jointly optimized with the diffusion loss \mathcal{L}_{DDPM} similar to Eq.6. We name such configuration as *Crossway-CURL*.

$$\begin{aligned}
 M_{sim} &= X_{t,a1}^k \cdot W \cdot sg(X_{t,a2,ema}^k)^T \\
 \mathcal{L}_{CURL} &= \text{CrossEntropyLoss}(M_{sim}, \text{range}(b)) \quad (7) \\
 \mathcal{L}_{\text{Crossway-CL}} &= \mathcal{L}_{DDPM} + \alpha \mathcal{L}_{CURL}
 \end{aligned}$$

From Tab. VII, the contrastive learning variant Crossway-CURL yields much worse performance compared to the baseline on multiple tasks, indicating that not all auxiliary SSL losses benefit policy learning. Such observations happen to align with the online reinforcement learning case [35].

V. RELATED WORK

A. Behavioral Cloning

Behavioral Cloning (BC) [1], [36], [37] is a straightforward but surprisingly effective way to obtain robot policies.

With pre-collected state-action pairs, BC learns a policy like fitting a dataset [38], with additional techniques like reward labeling / Inverse Reinforcement Learning (IRL) [39], [40], distribution matching [41], [42], and incorporating extra information [43], [44], [45]. Apart from explicitly generating output actions, BC can be done implicitly, where an energy-based model is learned to model the action distribution [33]. Implicit BC is found to be effective in real-world robot tasks. BC also boosts some online RL algorithms like TD3+BC [46], DeepMimic [47], and more [48]. Recent Diffusion-based BC is more like an advanced approach for matching the behavior distribution, which potentially helps mitigate the distribution shift problem in BC [49].

B. Policy Learning as Sequential Modeling

Sequential modeling [4], [5], [6] is the recent direction to solve offline-RL or Imitation Learning problems. The key is to optimize a policy on a trajectory basis from pre-collected experiences, with a reward signal (offline-RL) or without it (imitation learning). To model the trajectory composed by state-action-reward tuples, Transformer [9] is firstly adopted to this problem in the light of success in modeling natural language. In this formulation, state-action-reward tuples are regarded as equal units [4], [5] or with Markovian properties [6], [50] for better long-term modeling. There are works to extend the formulation to online learning [51], hindsight matching [52], and bootstrapping [53]. Recently, diffusion models have also been adopted to this problem [15], [17], [18], on which [54] conducts a comprehensive survey.

C. Self-supervised Learning

Self-supervised learning (SSL) [55] is used to learn data representations without task labels. SSL is commonly used to (pre-)train task-agnostic foundation models [56], [57], [58], [59], [60], [61], or is used as an auxiliary task together with other learning paradigms [35], [62], [63], [64]. Similarly, SSL has multiple ways to combine with policy learning, and we briefly categorize them into two: pre-training with SSL [65], [66], [67], [68], [69], [70], [71], [72], and policy learning jointly with auxiliary SSL tasks [73], [74], [75], [76], [77], [78], [34], [79], [80], [81], [82], [83], [84], [85], [86], [87], [35]. Studies [35] have shown that different ways to combine policy learning with SSL have different outcomes. In this work, we follow the joint learning style that optimizes diffusion and reconstruction objectives together.

VI. CONCLUSION

We explore the potential improvements for diffusion-based visuomotor policy when integrating an auxiliary self-supervised objective with the existing diffusion loss during training. Our proposed method, Crossway Diffusion, achieves consistent improvements across various challenging tasks, including two real-world scenarios, with a simple state decoder and an auxiliary reconstruction objective. Given the simplicity of our method and the detailed ablations, we hope this work will inspire further exploration into leveraging evolving self-supervised learning techniques for more effective diffusion-based policies.

REFERENCES

- [1] D. A. Pomerleau, "Alvinn: An autonomous land vehicle in a neural network," in *Advances in Neural Information Processing Systems (NeurIPS)*, 1988.
- [2] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, T. Jackson, S. Jesmonth, N. J. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, I. Leal, K.-H. Lee, S. Levine, Y. Lu, U. Malla, D. Manjunath, I. Mordatch, O. Nachum, C. Parada, J. Peralta, E. Perez, K. Pertsch, J. Quiambao, K. Rao, M. Ryoo, G. Salazar, P. Sanketi, K. Sayed, J. Singh, S. Sontakke, A. Stone, C. Tan, H. Tran, V. Vanhoucke, S. Vega, Q. Vuong, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich, "Rt-1: Robotics transformer for real-world control at scale," *Robotics science and systems (RSS)*, 2023.
- [3] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, X. Chen, K. Chormanski, T. Ding, D. Driess, A. Dubey, C. Finn *et al.*, "Rt-2: Vision-language-action models transfer web knowledge to robotic control," *arXiv preprint arXiv:2307.15818*, 2023.
- [4] L. Chen, K. Lu, A. Rajeswaran, K. Lee, A. Grover, M. Laskin, P. Abbeel, A. Srinivas, and I. Mordatch, "Decision transformer: Reinforcement learning via sequence modeling," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 34, pp. 15 084–15 097, 2021.
- [5] M. Janner, Q. Li, and S. Levine, "Offline reinforcement learning as one big sequence modeling problem," in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2021.
- [6] J. Shang, X. Li, K. Kahatapitiya, Y.-C. Lee, and M. S. Ryoo, "Starformer: Transformer with state-action-reward representations for robot learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–16, 2022.
- [7] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [8] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020.
- [9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2017.
- [10] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, X. Chen, K. Chormanski, T. Ding, D. Driess, A. Dubey, C. Finn *et al.*, "Rt-2: Vision-language-action models transfer web knowledge to robotic control," *arXiv preprint arXiv:2307.15818*, 2023.
- [11] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *International Conference on Machine Learning*. PMLR, 2015, pp. 2256–2265.
- [12] J. Ho, A. Jain, and P. Abbeel, "Denosing diffusion probabilistic models," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 6840–6851, 2020.
- [13] J. Song, C. Meng, and S. Ermon, "Denosing diffusion implicit models," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.
- [14] R. Burgert, X. Li, A. Leite, K. Ranasinghe, and M. S. Ryoo, "Diffusion illusions: Hiding images in plain sight," *arXiv preprint arXiv:2312.03817*, 2023.
- [15] M. Janner, Y. Du, J. B. Tenenbaum, and S. Levine, "Planning with diffusion for flexible behavior synthesis," *Proceedings of the International Conference on Machine Learning (ICML)*, 2022.
- [16] Z. Wang, J. J. Hunt, and M. Zhou, "Diffusion policies as an expressive policy class for offline reinforcement learning," *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.
- [17] T. Pearce, T. Rashid, A. Kanervisto, D. Bignell, M. Sun, R. Georgescu, S. V. Macua, S. Z. Tan, I. Momennejad, K. Hofmann *et al.*, "Imitating human behaviour with diffusion models," *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.
- [18] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *Robotics science and systems (RSS)*, 2023.
- [19] A. Mandlkar, D. Xu, J. Wong, S. Nasiriany, C. Wang, R. Kulkarni, L. Fei-Fei, S. Savarese, Y. Zhu, and R. Martín-Martín, "What mat-
ters in learning from offline demonstrations for robot manipulation," *Conference on Robot Learning (CoRL)*, 2021.
- [20] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, "Hierarchical text-conditional image generation with clip latents," *arXiv preprint arXiv:2204.06125*, 2022.
- [21] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 10 684–10 695.
- [22] S. Gu, D. Chen, J. Bao, F. Wen, B. Zhang, D. Chen, L. Yuan, and B. Guo, "Vector quantized diffusion model for text-to-image synthesis," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 10 696–10 706.
- [23] R. Burgert, K. Ranasinghe, X. Li, and M. S. Ryoo, "Peekaboo: Text to image diffusion models are zero-shot segmentors," *arXiv preprint arXiv:2211.13224*, 2022.
- [24] H.-C. Wang, S.-F. Chen, and S.-H. Sun, "Diffusion model-augmented behavioral cloning," *arXiv preprint arXiv:2302.13335*, 2023.
- [25] C. Lu, P. J. Ball, and J. Parker-Holder, "Synthetic experience replay," *arXiv preprint arXiv:2303.06614*, 2023.
- [26] T. Yu, T. Xiao, A. Stone, J. Tompson, A. Brohan, S. Wang, J. Singh, C. Tan, J. Peralta, B. Ichter *et al.*, "Scaling robot learning with semantically imagined experience," *Robotics science and systems (RSS)*, 2023.
- [27] Y. Dai, M. Yang, B. Dai, H. Dai, O. Nachum, J. Tenenbaum, D. Schuurmans, and P. Abbeel, "Learning universal policies via text-guided video generation," *arXiv preprint arXiv:2302.00111*, 2023.
- [28] J. Shang, K. Kahatapitiya, X. Li, and M. S. Ryoo, "Starformer: Transformer with state-action-reward representations for visual reinforcement learning," in *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer, 2022, pp. 462–479.
- [29] E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville, "Film: Visual reasoning with a general conditioning layer," in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, vol. 32, 2018.
- [30] N. Watters, L. Matthey, C. P. Burgess, and A. Lerchner, "Spatial broadcast decoder: A simple architecture for learning disentangled representations in vaes," *arXiv preprint arXiv:1901.07017*, 2019.
- [31] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," in *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer, 2020, pp. 405–421.
- [32] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017.
- [33] P. Florence, C. Lynch, A. Zeng, O. A. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson, "Implicit behavioral cloning," in *Conference on Robot Learning (CoRL)*. PMLR, 2022, pp. 158–168.
- [34] M. Laskin, A. Srinivas, and P. Abbeel, "Curl: Contrastive unsupervised representations for reinforcement learning," in *Proceedings of the International Conference on Machine Learning (ICML)*. PMLR, 2020, pp. 5639–5650.
- [35] X. Li, J. Shang, S. Das, and M. Ryoo, "Does self-supervised learning really improve reinforcement learning from pixels?" in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 35, 2022, pp. 30 865–30 881.
- [36] L. P. Kaelbling, *Learning in embedded systems*. MIT press, 1993.
- [37] P. Kormushev, S. Calinon, and D. G. Caldwell, "Imitation learning of positional and force skills demonstrated via kinesthetic teaching and haptic input," *Advanced Robotics*, vol. 25, no. 5, pp. 581–603, 2011.
- [38] C. G. Atkeson and S. Schaal, "Robot learning from demonstration," in *Proceedings of the International Conference on Machine Learning (ICML)*, vol. 97, 1997, pp. 12–20.
- [39] A. Y. Ng and S. Russell, "Algorithms for inverse reinforcement learning," in *Proceedings of the International Conference on Machine Learning (ICML)*, vol. 1, 2000, p. 2.
- [40] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, J. Peters *et al.*, "An algorithmic perspective on imitation learning," *Foundations and Trends® in Robotics*, vol. 7, no. 1-2, pp. 1–179, 2018.
- [41] J. Ho and S. Ermon, "Generative adversarial imitation learning," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 29, 2016.
- [42] X. B. Peng, Z. Ma, P. Abbeel, S. Levine, and A. Kanazawa, "Amp: Adversarial motion priors for stylized physics-based character con-

- trol,” *ACM Transactions on Graphics (TOG)*, vol. 40, no. 4, pp. 1–20, 2021.
- [43] P. Florence, L. Manuelli, and R. Tedrake, “Self-supervised correspondence in visuomotor policy learning,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 492–499, 2019.
- [44] T. Zhang, Z. McCarthy, O. Jow, D. Lee, X. Chen, K. Goldberg, and P. Abbeel, “Deep imitation learning for complex manipulation tasks from virtual reality teleoperation,” in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 5628–5635.
- [45] R. Rahmatizadeh, P. Abolghasemi, L. Bölöni, and S. Levine, “Vision-based multi-task manipulation for inexpensive robots using end-to-end learning from demonstration,” in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 3758–3765.
- [46] S. Fujimoto and S. S. Gu, “A minimalist approach to offline reinforcement learning,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 34, pp. 20 132–20 145, 2021.
- [47] X. B. Peng, P. Abbeel, S. Levine, and M. Van de Panne, “Deepmimic: Example-guided deep reinforcement learning of physics-based character skills,” *ACM Transactions On Graphics (TOG)*, vol. 37, no. 4, pp. 1–14, 2018.
- [48] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, “Learning complex dexterous manipulation with deep reinforcement learning and demonstrations,” *Robotics science and systems (RSS)*, 2018.
- [49] P. Abbeel and A. Y. Ng, “Apprenticeship learning via inverse reinforcement learning,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2004, p. 1.
- [50] S. Hu, L. Shen, Y. Zhang, and D. Tao, “Graph decision transformer,” *arXiv preprint arXiv:2303.03747*, 2023.
- [51] Q. Zheng, A. Zhang, and A. Grover, “Online decision transformer,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2022.
- [52] H. Furuta, Y. Matsuo, and S. S. Gu, “Distributional decision transformer for hindsight information matching,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2022.
- [53] K. Wang, H. Zhao, X. Luo, K. Ren, W. Zhang, and D. Li, “Bootstrapped transformer for offline reinforcement learning,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [54] Z. Zhu, H. Zhao, H. He, Y. Zhong, S. Zhang, Y. Yu, and W. Zhang, “Diffusion models for reinforcement learning: A survey,” *arXiv preprint arXiv:2311.01223*, 2023.
- [55] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol, “Extracting and composing robust features with denoising autoencoders,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2008, pp. 1096–1103.
- [56] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, “Masked autoencoders are scalable vision learners,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 16 000–16 009.
- [57] Z. Tong, Y. Song, J. Wang, and L. Wang, “Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 35, pp. 10 078–10 093, 2022.
- [58] L. Wang, B. Huang, Z. Zhao, Z. Tong, Y. He, Y. Wang, Y. Wang, and Y. Qiao, “Videomae v2: Scaling video masked autoencoders with dual masking,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 14 549–14 560.
- [59] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, “Language models are few-shot learners,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 1877–1901, 2020.
- [60] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *Proceedings of the International Conference on Machine Learning (ICML)*. PMLR, 2020, pp. 1597–1607.
- [61] L. Zhou, H. Liu, J. Bae, J. He, D. Samaras, and P. Prasanna, “Self pre-training with masked autoencoders for medical image classification and segmentation,” in *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2023, pp. 1–6.
- [62] R. Burgert, J. Shang, X. Li, and M. Ryoo, “Neural neural textures make sim2real consistent,” *Conference on Robot Learning (CoRL)*, 2022.
- [63] H. Yu, F. Wang, G. Teodoro, F. Chen, X. Guo, J. M. Nickerson, and J. Kong, “Self-supervised semantic segmentation of retinal pigment epithelium cells in flatmount fluorescent microscopy images,” *Bioinformatics*, vol. 39, no. 4, p. btad191, 2023.
- [64] S. Das, T. Jain, D. Reilly, P. Balaji, S. Karmakar, S. Marjit, X. Li, A. Das, and M. S. Ryoo, “Limited data, unlimited potential: A study on vits augmented by masked autoencoders,” in *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2024, pp. 6878–6888.
- [65] D. Ha and J. Schmidhuber, “Recurrent world models facilitate policy evolution,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 31, 2018.
- [66] P. Sermanet, C. Lynch, Y. Chebotar, J. Hsu, E. Jang, S. Schaal, S. Levine, and G. Brain, “Time-contrastive networks: Self-supervised learning from video,” in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1134–1141.
- [67] A. Zhan, P. Zhao, L. Pinto, P. Abbeel, and M. Laskin, “A framework for efficient robotic manipulation,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [68] R. Shah and V. Kumar, “Rrl: Resnet as representation for reinforcement learning,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.
- [69] A. Stooke, K. Lee, P. Abbeel, and M. Laskin, “Decoupling representation learning from reinforcement learning,” in *Proceedings of the International Conference on Machine Learning (ICML)*. PMLR, 2021, pp. 9870–9879.
- [70] J. Shang and M. S. Ryoo, “Self-supervised disentangled representation learning for third-person imitation learning,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 214–221.
- [71] C. Wang, X. Luo, K. Ross, and D. Li, “Vrl3: A data-driven framework for visual deep reinforcement learning,” *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [72] T. Xiao, I. Radosavovic, T. Darrell, and J. Malik, “Masked visual pre-training for motor control,” *arXiv preprint arXiv:2203.06173*, 2022.
- [73] A. v. d. Oord, Y. Li, and O. Vinyals, “Representation learning with contrastive predictive coding,” *arXiv preprint arXiv:1807.03748*, 2018.
- [74] M. Igl, L. Zintgraf, T. A. Le, F. Wood, and S. Whiteson, “Deep variational reinforcement learning for pomdps,” in *Proceedings of the International Conference on Machine Learning (ICML)*. PMLR, 2018, pp. 2117–2126.
- [75] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson, “Learning latent dynamics for planning from pixels,” in *Proceedings of the International Conference on Machine Learning (ICML)*. PMLR, 2019, pp. 2555–2565.
- [76] P. Yingjun and H. Xinwen, “Learning representations in reinforcement learning: An information bottleneck approach,” *arXiv preprint arXiv:1911.05695*, 2019.
- [77] D. Yarats, A. Zhang, I. Kostrikov, B. Amos, J. Pineau, and R. Fergus, “Improving sample efficiency in model-free reinforcement learning from images,” in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2021, pp. 10 674–10 681.
- [78] A. X. Lee, A. Nagabandi, P. Abbeel, and S. Levine, “Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 741–752, 2020.
- [79] J. Zhu, Y. Xia, L. Wu, J. Deng, W. Zhou, T. Qin, T.-Y. Liu, and H. Li, “Masked contrastive representation learning for reinforcement learning,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 03, pp. 3421–3433, 2023.
- [80] B. Mazoure, R. Tachet des Combes, T. L. Doan, P. Bachman, and R. D. Hjelm, “Deep reinforcement and infomax learning,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 3686–3698, 2020.
- [81] K.-H. Lee, I. Fischer, A. Liu, Y. Guo, H. Lee, J. Canny, and S. Guadarrama, “Predictive information accelerates learning in rl,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 11 890–11 901, 2020.
- [82] Z. D. Guo, B. A. Pires, B. Piot, J.-B. Grill, F. Althé, R. Munos, and M. G. Azar, “Bootstrap latent-predictive representations for multitask reinforcement learning,” in *Proceedings of the International Conference on Machine Learning (ICML)*. PMLR, 2020, pp. 3875–3886.
- [83] M. Schwarzer, A. Anand, R. Goel, R. D. Hjelm, A. Courville, and P. Bachman, “Data-efficient reinforcement learning with momentum predictive representations,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.

- [84] A. Zhang, R. McAllister, R. Calandra, Y. Gal, and S. Levine, "Learning invariant representations for reinforcement learning without reconstruction," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.
- [85] T. Yu, C. Lan, W. Zeng, M. Feng, Z. Zhang, and Z. Chen, "Playvirtual: Augmenting cycle-consistent virtual trajectories for reinforcement learning," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 34, 2021.
- [86] T. D. Kulkarni, A. Gupta, C. Ionescu, S. Borgeaud, M. Reynolds, A. Zisserman, and V. Mnih, "Unsupervised learning of object keypoints for perception and control," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 32, 2019.
- [87] X. Wang, L. Lian, and S. X. Yu, "Unsupervised visual attention and invariance for reinforcement learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 6677–6687.
- [88] T. Salimans and J. Ho, "Progressive distillation for fast sampling of diffusion models," *Proceedings of the International Conference on Learning Representations (ICLR)*, 2022.
- [89] Y. Song, P. Dhariwal, M. Chen, and I. Sutskever, "Consistency models," in *Proceedings of the International Conference on Machine Learning (ICML)*, 2023.