

Extraneousness-Aware Imitation Learning

Ray Chen Zheng^{*4,1}, Kaizhe Hu^{*4,1}, Zhecheng Yuan^{4,1}, Boyuan Chen³, Huazhe Xu^{1,2,4}

Abstract—Visual imitation learning provides an effective framework to learn skills from demonstrations. However, the quality of the provided demonstrations usually significantly affects the ability of an agent to acquire desired skills. Therefore, the standard visual imitation learning assumes near-optimal demonstrations, which are expensive or sometimes prohibitive to collect. Previous works propose to learn from *noisy* demonstrations; however, the noise is usually assumed to follow a context-independent distribution such as a uniform or gaussian distribution. In this paper, we consider another crucial yet underexplored setting — imitation learning with task-irrelevant yet locally consistent segments in the demonstrations (e.g., wiping sweat while cutting potatoes in a cooking tutorial). We argue that such noise is common in real world data and term them as “extraneous” segments. To tackle this problem, we introduce Extraneousness-Aware Imitation Learning (EIL), a self-supervised approach that learns visuomotor policies from third-person demonstrations with extraneous subsequences. EIL learns action-conditioned observation embeddings in a self-supervised manner and retrieves task-relevant observations across visual demonstrations while excluding the extraneous ones. Experimental results show that EIL outperforms strong baselines and achieves comparable policies to those trained with perfect demonstration on both simulated and real-world robot control tasks. The project page can be found here: <https://sites.google.com/view/eil-website>.

I. INTRODUCTION

Imitation learning (IL) enables intelligent agents to acquire various skills from demonstrations [1], [2]; recent advances also extend IL to the visual domain [3], [4], [5], [6]. However, in contrast to how humans learn from demonstrations, artificial agents usually require “clean” data sampled from expert policies. Some recent literatures [7], [8], [9], [10] propose methods to perform imitation learning from noisy demonstrations. However, many of these methods are state-based, and are limited by their requirements such as additional labels or assumptions about the noise. Despite the effort, real-world data may contain extraneous segments that can hardly be defined or labelled. For example, when learning to cut potatoes from videos, humans can naturally ignore some of the demonstrators’ extraneous actions like wiping sweat in the halfway. This distinction raises the question naturally: how can we leverage the rich range of unannotated visual demonstrations for imitation learning without being hindered by their noise?

In this paper, we propose Extraneousness-Aware Imitation Learning (EIL) that enables agents to imitate from noisy video demonstrations with extraneous segments. Our

method allows agents to identify extraneous subsequences via self-supervised learning and selectively perform imitation from the task-relevant parts. Specifically, we train an action-conditioned encoder through temporal cycle-consistency (TCC) loss [11] to obtain the embeddings of each observation. In this way, the observations of similar progress in the demonstrations will gain similar embeddings. Then, we propose an Unsupervised Voting-based Alignment algorithm (UVA) to filter task-irrelevant frames across video clips. Finally, we introduce a few tasks to benchmark the performance of imitation learning from noisy data with extraneous sequences.

We evaluate our method on multiple visual-input robot control tasks in both simulator and the real-world. The experiment results suggest that the proposed encoder can produce embeddings useful for extraneousness detection. As a result, EIL outperforms various baselines and achieves comparable performance to those trained with perfect demonstrations.

Our contributions can be summarized as follows: 1) We propose a meaningful yet underexplored setting of visual imitation learning from demonstrations with extraneous segments 2) We introduce Extraneousness-Aware Imitation Learning (EIL) that learns selectively from the task-relevant parts by leveraging action-conditioned embeddings and alignment algorithms. 3) We introduce datasets with extraneous segments over several simulated or real-world tasks and demonstrate our method’s empirical effectiveness.

II. RELATED WORKS

A. Learning from Noisy Demonstration

Imitation learning [1], [2], [12], [13] includes behavior cloning [14], [15] which aims to copy the behaviors from the demonstration, and inverse reinforcement learning [16] that infers the reward function for learning policies. However, these methods usually assume access to expert demonstrations, which are hard to obtain in practice.

Recent works try to tackle the imitation learning problem when the demonstrations are noisy. However, through this line of research [17], [18], [19], [20], [8], [9], [10], [21], the vast majority of the works are done in the low-dimensional state space rather than the high-dimensional image space. Furthermore, it is common in previous works [21] to assume the noise is sampled from an *a priori* distribution. Methods designed specifically for such noise might fail completely when the noise violates the assumption. Recently, more attentions are drawn to learning from realistic visual demonstrations, e.g., Chen et al. propose to learn policies from “in-the-wild” videos [22]. While the method achieves impressive

*Denotes equal contribution.

¹Tsinghua University, ²Shanghai AI Lab, ³Massachusetts Institute of Technology, ⁴Shanghai Qi Zhi Institute

Contact: zhengrc19@mails.tsinghua.edu.cn.

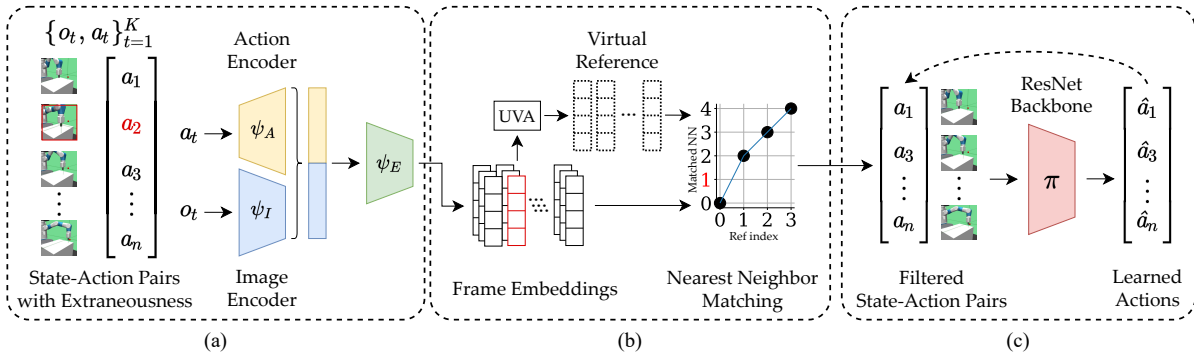


Fig. 1: **Extraneousness-Aware Imitation Learning (EIL)**. The overall framework contains 3 components: (a) It encodes the state-action pairs into representation. (b) It takes in the embeddings and process them with unsupervised voting-based alignment (UVA) algorithms. (c) It performs visual imitation learning with the aligned state action pairs. We note that (b) can be a simple filtering algorithm when reference trajectories are available.

results, they focus on dealing with diverse demonstrations without considering the “extraneousness” explicitly.

B. Self-Supervised Learning from Videos and its Application to Control and Robotics Tasks

Self-supervised learning (SSL) from videos can learn visual representations with temporal information for different downstream tasks from unlabeled data [23], [24], [25], [26], [27], [28], [29]. A recent line of research utilizes SSL for learning correspondences [30], [31], [32], [11], [33], [34]. Specifically, Dwibedi et al. propose to find correspondences across time in multiple videos with the help of cycle-consistency, where frames with similar progress will be encoded to similar embeddings [11], [35], [36]. These methods offer a welcoming way to leverage unlabeled and noisy real-world data. In recent years, SSL also promises to help with visuomotor tasks in control and robotics [37], [38]. For example, TCN [39] learns a self-supervised temporal-consistent embedding for imitation learning and reinforcement learning. XIRL [40] learns a self-supervised embedding that estimates task progress for inverse reinforcement learning. [41], [42], [43] directly map the observations such as images to the target domain. The distinction between EIL and previous works is that we tackle the problem where demonstrations have extraneous subsequences, rather than different visual appearances, view points, or embodiments.

III. METHOD

In this section, we first describe the problem setup and then introduce Extraneousness-Aware Imitation Learning (EIL), a simple yet effective approach for learning visuomotor policies from videos that have extraneous subsequences.

A. Problem Statement

We consider the setting where an agent aims at learning visuomotor policies from K video demonstrations $\{\mathcal{D}_i\}_{i=1}^K$. In the i^{th} video, the j^{th} observation o_j^i is paired up with its corresponding action a_j^i . For each sequence in the demonstration set, there are L extraneous subsequences $\{\mathcal{E}_n\}_{n=1}^L$ that are task-irrelevant yet locally consistent. In contrast to

existing works that have various assumptions about the noise, our setting only assumes each video to contain more than 50% of task-relevant content [44], [45].

The imitation agent takes a high-dimensional observation o_t as input and outputs an action \hat{a}_t at timestep t . To successfully imitate from the aforementioned demonstrations, the agent needs to reason about what the task-relevant parts are and rule out the extraneous subsequences.

B. Extraneousness-Aware Imitation Learning (EIL)

1) *Overview*: EIL is a general framework for imitating from videos with extraneous subsequences. The intuition of EIL is that the task-relevant parts among different demonstrations will share similar semantics meaning in the latent space, thus they can be aligned with each other. Following such intuition, when more than one demonstration sequence are given, we can match their embeddings to retrieve the task-relevant parts. In the case where a perfect reference demonstration is available, we can match frames in other sequences with that of the reference trajectory. However, in most cases, such a reference is hard to obtain. Hence, we propose an unsupervised alignment algorithm to retrieve task-relevant parts from a set of noisy demonstrations.

Figure 1 demonstrates the overview of EIL. In Figure 1(a), we learn a temporal representation of each frame conditioned on both its visual observation and action through temporal cycle consistency loss. After obtaining the representation, as shown in Figure 1(b), we propose an unsupervised voting method to perform video filtering when no perfect demonstration is available. Finally, as described in Figure 1(c), we perform standard visual imitation learning on top of the denoised data from the alignment procedure.

2) *Action-conditioned Temporal Cycle Consistency Representation Learning*: We first learn representations that encode temporal information for frame alignment across different video demonstrations. We train an image encoder ψ_I and an action encoder ψ_A that embed the observations and actions into corresponding features $\psi_I(o)$ and $\psi_A(a)$. Then, we concatenate $\psi_I(o)$ and $\psi_A(a)$ to a multi-layer perceptron (MLP) ψ_E to obtain the embeddings that have

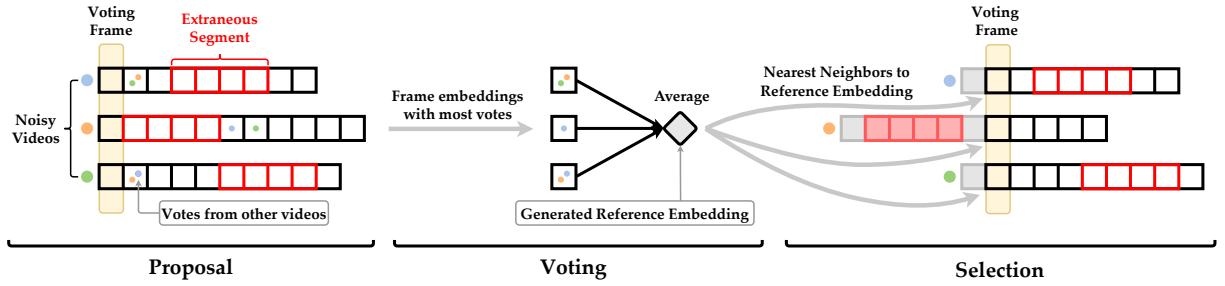


Fig. 2: **Unsupervised Voting-based Alignment (UVA)**. The three-stage structure includes: 1) Proposal where each frame performs nearest neighbor voting according to their embeddings. 2) Voting where all the selected frame embeddings are averaged out to get a virtual reference embedding. 3) Selection where we select the actual frame as our training data based on the estimated embedding.

temporal correspondences between two sequences. For simplicity, we use two demonstration sequences S and T and their computed embeddings $U = \{u_1, u_2, \dots, u_N\}$ and $V = \{v_1, v_2, \dots, v_M\}$ as an example. N and M denotes the sequence lengths respectively.

The main goal here is to encourage cycle-consistency between the two embedding sequences. For any $u_i \in U$, we find the nearest neighbor, $v_j = \arg \min_{v \in V} \|v - u_i\|$. Then we repeat the procedure and find $u_k = \arg \min_{u \in U} \|v_j - u\|$ which is the nearest neighbor for v_j . When $i = k$, the embedding u_i is cycle-consistent. To optimize the cycle-consistency, we use a differentiable matching loss: for the selected u_i , we compute the soft nearest neighbor by $\tilde{v} = \sum_j \alpha_j v_j$, where $\alpha_j = \frac{\exp(-\|u_i - v_j\|^2)}{\sum_k \exp(-\|u_i - v_k\|^2)}$. Then, we compute the “cycle-back” soft nearest neighbor to \tilde{v} similarly: $\tilde{u} = \sum_k \beta_k u_k$, where $\beta_k = \frac{\exp(-\|\tilde{v} - u_k\|^2)}{\sum_j \exp(-\|\tilde{v} - u_j\|^2)}$. The predicted index \hat{i} is calculated by $\hat{i} = \sum_k \beta_k k$. To obtain the cycle-consistency, \hat{i} should be close to the true i . Hence, we minimize the loss with an imposed Gaussian prior and variance regularization [11]

$$\mathcal{L} = \frac{|i - \hat{i}|^2}{\sigma^2} + \lambda \log(\sigma) \quad (1)$$

where $\sigma = \sum_k \beta_k (k - \hat{i})^2$ and λ is the regularization weight. At test time, the indices can be rounded to integers.

3) *Unsupervised Voting-based Frame Alignment (UVA)*: After obtaining the embedding for each observation-action pair, we try to align frames and drop extraneous frames according to a frame-wise similarity in the latent space.

To achieve this objective, we propose a voting-based frame matching algorithm that can remove the extraneous segments from a set of videos. A conceptual illustration is shown in Figure 2. For K demonstration sequences, we initially mark the first frame of each video as the “voting frame”. The distance and nearest neighbor mentioned below is in the embedding space. Our algorithm can be described as below:

- **Proposal.** For each video, find nearest neighbor to the “voting frame” among $K - 1$ other videos.

- **Voting.** In each video, the frame selected as nearest neighbor the most times is marked as a new “voting frame” of that video. We average the embeddings of all the newly selected “voting frames” to get a virtual reference embedding representing the current progress.
- **Selection.** We select the nearest neighbor of the virtual embedding in each video as the new “voting frame”. A causal restriction to select only in frames after the current “voting frame” is applied.

In a simpler setting where we have access to a perfect reference demonstration, our algorithm degenerates to simply picking frames in each video that are nearest neighbors to each frame of the reference.

4) *Visual Imitation Learning*: As shown in Figure 1(c), we perform standard visual imitation learning to learn a policy π that minimizes distance between the predicted actions and the ground-truth actions using the state-action pairs that are selected previously. Specifically, for continuous actions, we use the ℓ_2 loss: $L = \|a_i - \hat{a}_i\|^2$, where a_i and \hat{a}_i are the ground truth and predicted action respectively. For discrete actions, we use cross-entropy loss instead. We use ResNet-18 as our policy network to process the image input.

IV. EXPERIMENT

In this section, we describe our experiment setup and analyze the results. We compare our method with strong baselines on three simulated continuous control tasks. We also evaluate EIL on real-world robots. We aim to understand the extraneousness-aware imitation learning problem by answering the following questions: 1) Does EIL as a framework help the agent imitate from visual demonstrations that contain extraneousness? 2) Can the action-conditioned self-supervised representation differentiate between extraneous and task relevant components in the demonstrations? 3) What are the key factors and design choices for EIL?

A. Simulated Control Tasks

1) *Setup and Datasets*: Visualization and a brief description of the tasks are given in Figure 3. The *Reach* and *Push* tasks’ objectives are to move the gripper itself or an object to the target position. For *Stir*, the agent needs to place its gripper inside a bowl, then revolve it around the center of the

TABLE I: Averaged success rates and standard deviations of EIL and other baselines for simulated control tasks.

Task	Oracle	EIL (Ours)	Behavior Cloning	RL	TCN	RIL-Co	Random Policy
<i>Reach</i>	100%	85% ± 3%	77% ± 4%	20% ± 5%	80% ± 3%	82% ± 5%	4% ± 3%
<i>Push</i>	100%	79% ± 3%	66% ± 4%	16% ± 5%	76% ± 3%	22% ± 6%	0%
<i>Stir</i>	100%	83% ± 3%	55% ± 4%	32% ± 7%	61% ± 4%	54% ± 7%	0%

TABLE II: Averaged minimum distances and standard deviations of EIL and other baselines for simulated control tasks.

Task	Oracle	EIL (Ours)	Behavior Cloning	RL	TCN	RIL-Co	Random Policy
<i>Reach</i>	0.0175 ± 0.0105	0.0339 ± 0.0166	0.0374 ± 0.0174	0.0716 ± 0.0224	0.0371 ± 0.0172	0.0371 ± 0.0153	0.2059 ± 0.0675
<i>Push</i>	0.0242 ± 0.0060	0.0290 ± 0.0350	0.0403 ± 0.0435	0.1768 ± 0.1790	0.0305 ± 0.0275	0.1011 ± 0.0484	0.1456 ± 0.0152
<i>Stir</i>	0.0108 ± 0.0017	0.0388 ± 0.0173	0.0543 ± 0.0300	0.0626 ± 0.0045	0.0504 ± 0.0175	0.0723 ± 0.0410	1.2184 ± 0.1304




Fig. 3: **The simulated continuous control environments.** *Reach* (a) and *Push* (b) are goal conditioned environments where success is declared when target object is close enough to the destination. In *Stir* (c), success is declared if the end-effector correctly follows a target trajectory.

bowl. For each task, we collect two datasets – an extraneous dataset as well as a perfect dataset. In the perfect dataset, all the state-action pairs $\{s_i, a_i\}_{i=1\dots N}$ are sampled from an expert policy π_E . While in the extraneous dataset, there are one or more subsequences that contains $n-m+1$ consecutive steps $\{o_i, a_i\}_{i=m\dots n}$ sampled from another policy π_{ext} .

We note that we do not have any constraints on the non-expert policy π_{ext} , which means it can perform meaningful actions for an irrelevant task or totally random actions. The objects and targets are randomly chosen for every demonstration. For the extraneous parts, we insert locally consistent action sequences at random timesteps. Specifically, for *Reach* and *Push*, the extraneous action is the agent deviating away from its original trajectory at a random timestep then coming back. For *Stir*, the agent moves the gripper outside the bowl towards a random position and returns, which simulates a real-world scenario where a human fetches an object in the middle of stirring. In the test time, we either provide one perfect reference trajectory or no reference trajectory.

2) *Visual Imitation Learning*: For continuous control tasks, we optionally give our policy access to the intrinsic low-level state information of the agent (i.e., position and velocity of the gripper, but not the goal or object). We train a total of 3 random seeds, then choose the one with better performance in validation. We use the selection scheme from Husenot et al. for hyperparameter tuning. [46].

3) *Metrics and Evaluation*: For the goal-conditioned continuous control tasks (*Reach* and *Push*), we use both success rate and minimum goal-object or gripper-object distance as evaluation metrics. For *Stir*, we calculate the mean deviation between the end-effector’s trajectory and the target circular

orbit, and use this deviation to obtain success rate. All the experiments are conducted with 3 random seeds. We average over 50 trials for each seed and report the mean and standard deviation. For all the continuous control tasks, we use the default threshold value defined in the *Reach* and *Push* environment of 0.05 m.

4) *Baselines*: We compare EIL with several baselines.

Behavior cloning (BC). We train a neural network to imitate the mapping from observations to actions through vanilla supervised learning.

Reinforcement learning with embedding-based reward. We use reinforcement learning to accomplish the task. Instead of the sparse reward of whether the goal is reached, we first obtain a “goal embedding” by averaging over the emddings of each video’s last frame, then set a dense reward of the negative ℓ_2 distance between the current state-action embedding and the goal embedding. We note that this baseline has the privilege of interacting with the environment.

Time-contrastive networks (TCN) [39]. TCN is the predecessor model of TCC, and is originally designed for imitation and reinforcement learning on synchronized multi-view video demonstrations. In our paper, we adopt TCN by substituting the TCC encoder with a TCN network, while the rest of our method (i.e., UVA and visual IL) is unchanged.

Robust imitation learning with co-pseudo-labeling (RIL-Co) [7] RIL-Co is the state-of-the-art method for imitation learning from noisy demonstrations. We provide low-level state and action pairs as well as access to the environment to train the policy of RIL-Co. We note that 1) RIL-Co is not a visual imitation learning method and has the advantage of learning from low-level states, and 2) like the RL baseline, RIL-Co requires interaction with the environment.

Random policy. We sample actions from a uniform random distribution as baseline.

5) *Experimental Results*:

a) *Control Results*: Table I summarizes the mean and standard deviation of success rate on all tasks. The performance of vanilla behavior cloning degrades heavily when extraneousness is present in training data.

Reinforcement learning algorithms struggle to master any of the tasks, even with access to the environment. This may be due to the extraneous frames that cause the task-relevant embeddings to no longer linearly approach the goal embedding. With the continuous tendency of the embedding

sequence broken, the distance to the goal embedding may give a wrong guidance, causing the agent to fail. TCN shows good results in *Reach* and *Push*, but performs poorly in *Stir*. *Stir* is different from the other two tasks as it is periodic and could repeat the same states and actions. This violates the assumption of TCN: two frames that are distant in time should have distinct embeddings. RIL-Co, despite its privilege of having access to the environment and learning directly from low-level states, can only achieve comparable results in the simplest *Reach* task and cannot perform well in other tasks. A possible reason is the extraneous subsequences violates RIL-Co’s assumption for the noise part to be random and inconsistent. Our method outperforms all the baseline methods and demonstrates its ability to learn from extraneous-rich demonstrations. We also find that EIL outperforms other methods in terms of minimum distance as shown in Table II.




Fig. 4: **UVA’s alignment results.** We mark the extraneous subsequence with red shade, observations chosen by UVA with continuous lines, and observations filtered out with dashed lines. Our method successfully skips extraneous part while keeps the others.

b) *UVA Results:* To evaluate the ability for UVA to exclude task irrelevant frames, we visualize the filter result for each of the task in Figure 4. Despite some occasional confusion at the border area, UVA can successfully ignore the extraneous parts in the demonstrations. Detailed filtering results could be found at Table III and V.

B. Real-World Robot Control Tasks

To further evaluate the performance of EIL and demonstrate the practical value of our method, we adopt a real-world robot arm to learn from demonstrations with extraneous-ness. We train the arm on the tasks of *Reach* and *Push* using an extraneous dataset with the methods of EIL and behavior cloning.

1) *Setup and Datasets:* We use a Franka Emika Panda robot arm as shown in Figure 5. The RGB images are captured by an Intel RealSense D435i depth camera.

Similar to the simulator, we collect an extraneous dataset and a perfect dataset. The perfect dataset contains state-action pairs sampled purely from the export policy π_E , while the extraneous dataset contains one or more subsequences of state-action pairs sampled from another policy π_{ext} . In practice, the camera takes a picture of the current arm, object (in *Push*), and target. A human expert then controls the arm to take the correct action, generating a state-action pair and leading to the next state. Due to this process’s costly




Fig. 5: **Real-World Robot Setup.** The RealSense D435i (a) captures RGB images. The target for *Reach* (b) is a paper cup, while *Push* (c) uses a cylinder as its object and white circle as target.

nature, less demonstrations (20 videos of 70–100 frames) are collected, and the action space is changed from continuous to discrete (front, back, up, down, left, right).

2) *Visual Imitation Learning and Evaluation:* The visual imitation learning process of the real-world robot arm is identical to that of the simulator. However, data augmentation methods of random resizing and cropping are used to overcome jitters of camera positions and illumination conditions.

For evaluation, we measure the distances between the target and the gripper (*Reach*) or object (*Push*). There is also a timestep limit set for every evaluation trajectory. If the agent cannot reach a success state within the time limit, the trajectory will not be considered as successful. Distances are measured from the gripper or object to the center of the target, and the threshold distance of success is the diameter of the target. We evaluate success rate over 20 trials.

TABLE III: Extraneous Percentage (lower is better) and Success Rates (higher is better) for Real-World Robot Learning.

Task	Metric	BC	EIL (Ours)
<i>Reach</i>	Extraneous %	34.6%	9.6%
	Success Rate	40% ± 11%	80% ± 9%
<i>Push</i>	Extraneous %	39.4%	4.1%
	Success Rate	10% ± 7%	75% ± 10%

3) *Experimental Results:* Table III give the results of our real-world robot arm learning experiment. Our method significantly reduces the extraneous ratio in both datasets and improves success rate for imitation learning.

In *Reach*, the BC agent often gets stuck in a local area due to the extraneous segments. For *Push*, failures of BC are mainly due to two reasons: the gripper goes down too much that it hits the table surface, or the object slides out of the gripper’s control. The EIL agent, on the other hand, overcomes these difficulties smoothly with the help of UVA.

C. Visualized Results

We visualize an extraneousness-rich demonstration trajectory and the extracted subsequences with EIL in Figure 6. From the visualization, we see that the extraneous subsequence is filtered out, and the new demonstration has higher quality, leading to better train results.




Fig. 6: **Visualized result of filtered demonstration.** The extraneous contents (heading to the other corner in the mid-way) in *Push* are successfully filtered out by UVA. Gripper tips are marked by yellow dots in this graph for clarity.

D. Ablation Study

We conduct ablation study on our representation learning and the unsupervised alignment parts of EIL.

1) *Representation Learning*: EIL learns action-conditioned embeddings from the image encoder ψ_I and action encoder ψ_A . For continuous control tasks, intrinsic states are also available since the agent has information of itself. Therefore, we may add an intrinsic state encoder ψ_S in addition to ψ_I and ψ_A to provide this extra information. Here we ablate the effect of adding intrinsic state on task success rate and the quality of extraneous frame filtering.

a) *Effects on success rates*: As shown in Table IV, EIL is able to obtain the highest success rate in most of the cases. Meanwhile, adding ψ_S does not help most of the times. Table IV also infers that adding states during IL training does not guarantee improvement in final performance.

TABLE IV: Success rates for different encoder configurations to obtain temporal embeddings. Higher is better.

Task	IL trained	BC	EIL(ψ_I)	EIL($\psi_{I,A}$)	EIL($\psi_{I,A,S}$)	Oracle
<i>Reach</i>	w/ state	77%	85%	85%	80%	99%
	w/o state	74%	70%	78%	88%	100%
<i>Push</i>	w/ state	66%	67%	77%	74%	97%
	w/o state	57%	76%	79%	73%	100%
<i>Stir</i>	w/ state	55%	47%	74%	75%	100%
	w/o state	54%	38%	83%	72%	100%

b) *Effects on extraneous filtration*: Table V shows the percentage of extraneous content in the datasets before and after filtering. UVA is able to decrease extraneous percentage from over 25% to around 5% in every task.

TABLE V: Percentage of extraneous content before and after filtering by different encoder configurations. Lower is better.

Task	Quantity	Original	EIL(ψ_I)	EIL($\psi_{I,A}$)	EIL($\psi_{I,A,S}$)
<i>Reach</i>	Extraneous	624	108	137	91
	Total	2200	1650	2001	1299
	Extraneous %	28.4%	6.5%	6.8%	7.0%
<i>Push</i>	Extraneous	624	58	63	144
	Total	2489	2245	2245	2323
	Extraneous %	25.1%	2.6%	2.8%	6.2%
<i>Stir</i>	Extraneous	624	34	48	17
	Total	2040	1221	1338	987
	Extraneous %	30.6%	2.8%	3.6%	1.7%

2) *UVA*: In this paper, we mainly consider the setting where all the data are not perfect. However, in certain cases, a small amount of perfect demonstrations are also available. We compare the performance of EIL with and without perfect

reference trajectory in Table VI. We use two alignment methods with perfect reference demonstrations: dynamic time warping (DTW) [47], and nearest neighbor matching (NN). Nearest neighbor matching directly maps every frame of the perfect demonstration to its nearest neighbor in other videos. DTW ensures a smoother match curve that is chronological. We notice that UVA greatly outperforms DTW and NN despite not having access to perfect data. DTW performs even worse than vanilla Nearest Neighbor (NN) since the locally consistent noise segment disturbs the algorithm and often makes it stick to the segment once it steps into the extraneous area.

TABLE VI: Percentage of extraneous content in demonstrations before and after filtering, with or without perfect reference demonstrations. Lower is better.

Task	Quantity	Original	UVA	DTW	NN
<i>Reach</i>	Extraneous	624	137	237	152
	Total	2200	2001	2040	2040
	Extraneous %	28.4%	6.8%	11.6%	7.5%
<i>Push</i>	Extraneous	624	63	180	98
	Total	2489	2245	2440	2440
	Extraneous %	25.1%	2.8%	7.4%	4.0%
<i>Stir</i>	Extraneous	624	48	357	195
	Total	2040	1338	2040	2040
	Extraneous %	30.6%	3.6%	17.5%	9.6%

V. CONCLUSIONS AND FUTURE WORK

This paper focuses on visual imitation from demonstrations where temporally consistent yet task-irrelevant subsequences are present. We propose Extraneousness-Aware Imitation Learning (EIL), a framework that enables agents to identify extraneous subsequences from visual demonstrations via self-supervised learning. Empirical results show that EIL outperforms strong baselines on continuous and discrete control benchmarks in both simulator and the real-world. An exciting direction for future work is to integrate EIL with robotic learning from human demonstration methods, which usually deals with videos containing rich task-irrelevant segments. Another extension would be scenarios where observations are vulnerable to temporary corruption, like the dazzling light in a driving scene. EIL could filter out these corrupted observations as well since they are inconsistent with the normal observation sequence.

REFERENCES

- [1] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and autonomous systems*, vol. 57, no. 5, pp. 469–483, 2009.
- [2] S. Schaal *et al.*, "Learning from demonstration," *Advances in neural information processing systems*, pp. 1040–1046, 1997.
- [3] D. Pathak, P. Mahmoodeh, G. Luo, P. Agrawal, D. Chen, Y. Shentu, E. Shelhamer, J. Malik, A. A. Efros, and T. Darrell, "Zero-shot visual imitation," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2018, pp. 2050–2053.
- [4] 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 *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 5628–5635.
- [5] S. Young, D. Gandhi, S. Tulsiani, A. Gupta, P. Abbeel, and L. Pinto, "Visual imitation made easy," *arXiv preprint arXiv:2008.04899*, 2020.
- [6] Y. Zhu, Z. Wang, J. Merel, A. Rusu, T. Erez, S. Cabi, S. Tunyasuvunakool, J. Kramár, R. Hadsell, N. de Freitas, *et al.*, "Reinforcement and imitation learning for diverse visuomotor skills," *arXiv preprint arXiv:1802.09564*, 2018.
- [7] V. Tangkaratt, N. Charoenphakdee, and M. Sugiyama, "Robust imitation learning from noisy demonstrations," *arXiv preprint arXiv:2010.10181*, 2020.
- [8] D. Brown, W. Goo, P. Nagarajan, and S. Niekum, "Extrapolating beyond suboptimal demonstrations via inverse reinforcement learning from observations," in *International conference on machine learning*. PMLR, 2019, pp. 783–792.
- [9] Y.-H. Wu, N. Charoenphakdee, H. Bao, V. Tangkaratt, and M. Sugiyama, "Imitation learning from imperfect demonstration," in *International Conference on Machine Learning*. PMLR, 2019, pp. 6818–6827.
- [10] V. Tangkaratt, B. Han, M. E. Khan, and M. Sugiyama, "Variational imitation learning with diverse-quality demonstrations," in *International Conference on Machine Learning*. PMLR, 2020, pp. 9407–9417.
- [11] D. Dwibedi, Y. Aytar, J. Tompson, P. Sermanet, and A. Zisserman, "Temporal cycle-consistency learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1801–1810.
- [12] J. Ho and S. Ermon, "Generative adversarial imitation learning," *Advances in neural information processing systems*, vol. 29, pp. 4565–4573, 2016.
- [13] R. Dadashi, L. Hussenot, M. Geist, and O. Pietquin, "Primal wasserstein imitation learning," *arXiv preprint arXiv:2006.04678*, 2020.
- [14] M. Bain and C. Sammut, "A framework for behavioural cloning," in *Machine Intelligence 15*, 1995, pp. 103–129.
- [15] D. A. Pomerleau, "Efficient training of artificial neural networks for autonomous navigation," *Neural computation*, vol. 3, no. 1, pp. 88–97, 1991.
- [16] A. Y. Ng, S. J. Russell, *et al.*, "Algorithms for inverse reinforcement learning," in *ICML*, vol. 1, 2000, p. 2.
- [17] M. Kaiser, H. Friedrich, and R. Dillmann, "Obtaining good performance from a bad teacher," in *Programming by Demonstration vs. Learning from Examples Workshop at ML*, vol. 95. Citeseer, 1995.
- [18] D. H. Grollman and A. G. Billard, "Robot learning from failed demonstrations," *International Journal of Social Robotics*, vol. 4, no. 4, pp. 331–342, 2012.
- [19] B. Kim, A.-m. Farahmand, J. Pineau, and D. Precup, "Learning from limited demonstrations," in *NIPS*. Citeseer, 2013, pp. 2859–2867.
- [20] B. Burchfiel, C. Tomasi, and R. Parr, "Distance minimization for reward learning from scored trajectories," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 30, no. 1, 2016.
- [21] F. Sasaki and R. Yamashina, "Behavioral cloning from noisy demonstrations," in *International Conference on Learning Representations*, 2020.
- [22] A. S. Chen, S. Nair, and C. Finn, "Learning generalizable robotic reward functions from "in-the-wild" human videos," *arXiv preprint arXiv:2103.16817*, 2021.
- [23] D. Gordon, K. Ehsani, D. Fox, and A. Farhadi, "Watching the world go by: Representation learning from unlabeled videos," *arXiv preprint arXiv:2003.07990*, 2020.
- [24] D. Pathak, R. Girshick, P. Dollár, T. Darrell, and B. Hariharan, "Learning features by watching objects move," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 2701–2710.
- [25] N. Srivastava, E. Mansimov, and R. Salakhudinov, "Unsupervised learning of video representations using lstms," in *International conference on machine learning*. PMLR, 2015, pp. 843–852.
- [26] M. Mathieu, C. Couprie, and Y. LeCun, "Deep multi-scale video prediction beyond mean square error," *arXiv preprint arXiv:1511.05440*, 2015.
- [27] D. Jayaraman and K. Grauman, "Learning image representations tied to ego-motion," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 1413–1421.
- [28] P. Agrawal, J. Carreira, and J. Malik, "Learning to see by moving," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 37–45.
- [29] R. Goroshin, J. Bruna, J. Tompson, D. Eigen, and Y. LeCun, "Unsupervised learning of spatiotemporally coherent metrics," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 4086–4093.
- [30] A. Jabri, A. Owens, and A. A. Efros, "Space-time correspondence as a contrastive random walk," *arXiv preprint arXiv:2006.14613*, 2020.
- [31] X. Wang, A. Jabri, and A. A. Efros, "Learning correspondence from the cycle-consistency of time," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2566–2576.
- [32] C. Vondrick, A. Shrivastava, A. Fathi, S. Guadarrama, and K. Murphy, "Tracking emerges by coloring videos," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 391–408.
- [33] I. Hadji, K. G. Derpanis, and A. D. Jepson, "Representation learning via global temporal alignment and cycle-consistency," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 11 068–11 077.
- [34] S. Purushwalkam, T. Ye, S. Gupta, and A. Gupta, "Aligning videos in space and time," in *European Conference on Computer Vision*. Springer, 2020, pp. 262–278.
- [35] T. Zhou, P. Krahenbuhl, M. Aubry, Q. Huang, and A. A. Efros, "Learning dense correspondence via 3d-guided cycle consistency," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 117–126.
- [36] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232.
- [37] S. Nair, A. Rajeswaran, V. Kumar, C. Finn, and A. Gupta, "R3m: A universal visual representation for robot manipulation," *arXiv preprint arXiv:2203.12601*, 2022.
- [38] T. Xiao, I. Radosavovic, T. Darrell, and J. Malik, "Masked visual pre-training for motor control," *arXiv preprint arXiv:2203.06173*, 2022.
- [39] 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 *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 1134–1141.
- [40] K. Zakka, A. Zeng, P. Florence, J. Tompson, J. Bohg, and D. Dwibedi, "Xirl: Cross-embodiment inverse reinforcement learning," *arXiv preprint arXiv:2106.03911*, 2021.
- [41] Q. Zhang, T. Xiao, A. A. Efros, L. Pinto, and X. Wang, "Learning cross-domain correspondence for control with dynamics cycle-consistency," *arXiv preprint arXiv:2012.09811*, 2020.
- [42] L. Smith, N. Dhawan, M. Zhang, P. Abbeel, and S. Levine, "Avid: Learning multi-stage tasks via pixel-level translation of human videos," *arXiv preprint arXiv:1912.04443*, 2019.
- [43] H. Xiong, Q. Li, Y.-C. Chen, H. Bharadhwaj, S. Sinha, and A. Garg, "Learning by watching: Physical imitation of manipulation skills from human videos," *arXiv preprint arXiv:2101.07241*, 2021.
- [44] D. Angluin and P. Laird, "Learning from noisy examples," *Machine Learning*, vol. 2, no. 4, pp. 343–370, 1988.
- [45] N. Natarajan, I. S. Dhillon, P. K. Ravikumar, and A. Tewari, "Learning with noisy labels," *Advances in neural information processing systems*, vol. 26, pp. 1196–1204, 2013.
- [46] L. Hussenot, M. Andrychowicz, D. Vincent, R. Dadashi, A. Raichuk, S. Ramos, N. Momchev, S. Girgin, R. Marinier, L. Stafiniak, *et al.*, "Hyperparameter selection for imitation learning," in *International Conference on Machine Learning*. PMLR, 2021, pp. 4511–4522.
- [47] R. Bellman and R. Kalaba, "On adaptive control processes," *IRE Transactions on Automatic Control*, vol. 4, no. 2, pp. 1–9, 1959.