
AVID: Learning Multi-Stage Tasks via Pixel-Level Translation of Human Videos

Laura Smith Nikita Dhawan Marvin Zhang Pieter Abbeel Sergey Levine

Abstract

Reinforcement learning (RL) holds the promise of enabling robots to learn complex behaviors through experience; however, realizing this promise for long-horizon tasks in the real world requires mechanisms to reduce human burden in terms of defining the task and scaffolding the learning process. We study how these challenges can be alleviated with an automated robotic learning framework, in which multi-stage tasks are defined simply by providing videos of a human demonstrator and then learned by the robot from raw images. A central challenge in imitating human videos is the difference in appearance between the human and robot, which we reconcile by performing pixel-level image translation via CycleGAN to convert the human video into that of a robot, which can then be used to construct a reward function for RL. The robot then learns the task one stage at a time while occasionally asking the human to provide sparse, binary feedback in favor of human-provided resets at every stage of learning. This makes the learning process largely automatic, from intuitive task specification via a video to automated training with minimal human intervention.

1. Introduction

Consider a robot learning to operate a coffee machine directly from visual inputs. This task is representative of many tasks we would like robots to learn in the real world in that, though the overall skill is complex, it can be broken down into multiple stages, each of which is simple to learn. In principle, RL methods enable robots to learn skills such as this through their own experience. However, taking advantage of the natural stage-wise decomposition of the task is difficult with RL due to the challenges associated with defining a reward function for each stage, which often requires substantial human effort.

We may instead consider approaching this problem with imitation learning (IL), as utilizing demonstrations is an effective way to convey intricate goals and to circumvent exploration challenges (Argall et al., 2009). One of the

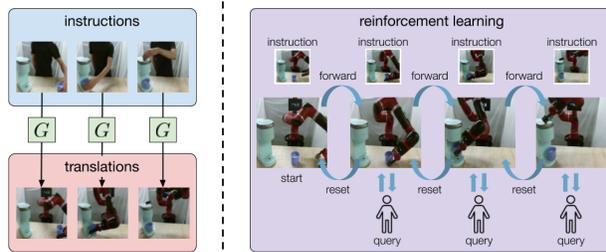


Figure 1. Left: Human instructions for each stage (top) are translated at the pixel level into robot instructions (bottom) via CycleGAN. Right: The robot attempts the task stage-wise, automatically resetting and retrying until the instruction classifier signals success, prompting the human to confirm via key press.

primary challenges with this approach is that robotic IL typically relies on demonstrations provided on the robot itself. The means by which these demonstrations can be obtained are often laborious and limited in applicability (Akgun et al., 2012; Calinon et al., 2009; Pastor et al., 2011; Zhang et al., 2018). For instance, kinesthetic teaching requires a person to guide the robot by hand, which is not only unintuitive and physically demanding, but also introduces distractions or obstructions in image-based learning. Teleoperation, another common method for providing demonstrations, can require specialized hardware and generally demands considerable operator expertise.

Humans learn to imitate in a starkly different manner: through observation. In contrast to being remote controlled or physically maneuvered, humans can learn multi-stage tasks by simply *watching* other people, imagining how they would perform the task themselves, and then practicing each stage of the task on their own. The question we ask in this work is: can we endow robots with the same ability?

We answer this question through the setting of learning from videos of human demonstrations. This setting eases the burden associated with providing demonstrations directly on the robot, as human demonstrations require comparatively minimal setup, hardware, and expertise. Thus, videos are often substantially easier for the user to provide, while opening up a much wider source of supervision for robots to learn skills. We still wish to leverage the demonstrations to provide natural, fine-grained guidance,

such that the robot can learn effectively by exploiting the stage-wise nature of the task. However, as we learn from visual inputs, differences in appearance between the human and robot must be reconciled.

To tackle this challenge, we propose a robotic learning system that adopts an approach similar to the strategy described above: one of imagination and practice. Starting from a human demonstration, our method *translates* this demonstration, at the pixel level, into images of the robot performing the task (see Figure 1). To remove the need for defining manual correspondences between the human and robot, we learn this translation via CycleGAN (Zhu et al., 2017), a framework for learning unpaired image-to-image translation. In order to handle multi-stage tasks, we reduce the human demonstration to a few *instruction images*, denoting the stages of the task. These instruction images are translated using CycleGAN, which are then used to provide a reward for model-based RL, enabling the robot to practice the skill to learn its physical execution. This phase is largely automated, as the robot learns to reset each stage on its own to practice it multiple times. Human supervision is only needed in the form of selectively requested key presses, indicating success or failure, and a few manual resets. We demonstrate that this approach is capable of solving complex, long-horizon tasks with minimal human involvement, removing most of the human burden associated with instrumenting the task setup, manually resetting the environment, and supervising the learning process.

We name our method automated visual instruction-following with demonstrations (AVID), and this method is the main contribution of our work. We evaluate AVID on two tasks using a Sawyer robot arm: operating a coffee machine and retrieving a cup from a drawer. AVID outperforms ablations and prior methods in terms of data efficiency and task success, learning coffee making with only 30 human demonstrations, amounting to 20 minutes of human demonstration time, and 180 minutes of robot interaction time with the environment.

2. Related Work

RL for Robotics. RL is a powerful framework for enabling robots to autonomously learn a wide range of real-world skills (Kalashnikov et al., 2018; Kober & Peters, 2011; Kober et al., 2013; Kohl & Stone, 2004; OpenAI et al., 2018; van Hoof et al., 2015). However, reward specification remains a major obstacle in deploying RL agents outside of controlled settings, as measuring rewards typically requires highly instrumented, and sometimes creative, setups such as motion capture and object tracking (Kormushev et al., 2010; Rusu et al., 2017), accelerometers to detect door opening (Yahya et al., 2017), or thermal cameras to detect pouring (Schenck & Fox, 2017). Our work, in

line with IL (Argall et al., 2009), circumvents the need for a reward function by using demonstrations.

For robotic applications, data efficiency is often another significant issue. A popular approach to improving efficiency is to use state representation learning to reduce the dimensionality of the data used for RL (Lesort et al., 2018). Another is to use model-based RL, which typically has been more efficient than model-free RL methods in practice (Chua et al., 2018; Deisenroth et al., 2014). We frame representation learning and model learning in our work as a joint variational inference procedure, combining aspects of prior work to model real world images and sparse rewards (Fu et al., 2018; Lee et al., 2019a). However, we consider multi-stage tasks that are significantly more complex than those studied in these prior works.

Learning multi-stage tasks. A common paradigm for learning multi-stage tasks is to learn or plan separately for each stage, through methods such as chaining dynamic movement primitives (Stulp et al., 2012) or linear-Gaussian controllers (Han et al., 2015) or learning inverse dynamics models (Nair et al., 2017; Sharma et al., 2019). These works either assume access to reward functions for each stage or demonstrations consistent with the robot’s observations and actions. A complementary line of work studies automatically segmenting demonstrations into stages or subgoals using techniques such as movement primitive libraries (Manschitz et al., 2014) or generative models (Kurutach et al., 2018). We assume that the task stages are easy to manually specify from the human demonstrations, so we leave automatic stage segmentation for future work. Instead, we focus on reducing human burden during the learning process.

One way in which we reduce human burden is by enabling the robot to reset each stage of the task, rather than requiring constant manual resets. This is similar to (Han et al., 2015), which proposes an algorithm that uses learned resets and can solve a multi-stage toy wrench task. However, we evaluate on image-based tasks and we do not assume access to any reward functions. Thus, there is no mechanism for automatically testing whether a stage has been completed, and instead we use human key presses to indicate the success or failure of a stage. Using human feedback has been studied in several recent works (Christiano et al., 2017; Daniel et al., 2014; Singh et al., 2019), and we further decrease the data and human supervision requirements compared to these works, in order to learn multi-stage tasks in only a few hours of autonomous robot execution, a few manual resets, and less than 150 key presses.

Learning from human demonstrations. As mentioned earlier, prior methods for robotic IL have typically used demonstrations consisting of observations and actions from the robot’s observation and action space (Argall et al.,

2009; Billard & Mataric, 2001; Hayes & Demiris, 1994; Kalakrishnan et al., 2013; Kober & Peters, 2009; Schaal et al., 2003). We relax standard assumptions that the expert data are given using the robot’s own embodiment. Some works have studied this setting using explicit pose and object detection (Kuniyoshi et al., 1994; Lee et al., 2013; Miyamoto et al., 1996; Ramirez-Amaro et al., 2017; Yang et al., 2015), essentially resolving the correspondence problem (Nehaniv & Dautenhahn, 2002) by instrumenting paired data collection or manually matching hand-specified key points. Other approaches have included predictive modeling (Rhinehart & Kitani, 2017; Tow et al., 2017), context translation (Liu et al., 2018; Sharma et al., 2019), learning reward representations (Sermanet et al., 2017; 2018), and meta-learning (Yu et al., 2018). In contrast, we explicitly account for the change in embodiment through learned pixel-level translation, we evaluate on long-horizon multi-stage tasks, and we do not assume access to any demonstrations given directly on the robot.

We evaluate the single-view version of time-contrastive networks (TCN) (Sermanet et al., 2018) on our tasks in Section 5. This prior method also handles embodiment changes, evaluates on visual robotic tasks, does not require demonstrations given on the robot, and has successfully learned pouring from videos of human pouring. TCN typically requires multiple views of human demonstrations, but we do not assume access to these in our problem setting, thus we compare to single-view TCN.

3. Preliminaries

3.1. Unsupervised Image-to-Image Translation

We approach human to robot demonstration translation as an unsupervised image-to-image translation problem, where the goal is to map images from a source domain X (human images, in our case) to a target domain Y (robot images) in the absence of paired training data (Liu et al., 2017; Zhu et al., 2017). Although the training data is not paired, we assume there exists a bijective mapping that can be uncovered in an unsupervised manner. This assumption is manifested in our experiments in Section 5 as similarities in morphology and perspective between human and robot images. The approach we use to learn this translation is CycleGAN (Zhu et al., 2017), a framework which learns two mappings: $G : X \rightarrow Y$ translates source to target, and $F : Y \rightarrow X$ translates target to source. The translations are learned in order to fool discriminators D_Y and D_X , which are trained to distinguish real images from translations in the target and source domains, respectively. This leads to a loss of the form

$$\mathcal{L}_{\text{GAN}}(G, D_Y) = \mathbb{E}[\log D_Y(y) + \log(1 - D_Y(G(x)))] ,$$

and similarly for F and D_X , where x and y are source and target images drawn from the data distribution. An additional loss promotes cycle consistency between G and F :

$$\mathcal{L}_{\text{cyc}}(G, F) = \mathbb{E}[\|x - F(G(x))\|_1 + \|y - G(F(y))\|_1] .$$

The overall training objective for CycleGAN is given by

$$\mathcal{L}_{\text{CG}}(G, F, D_X, D_Y) = \mathcal{L}_{\text{GAN}}(G, D_Y) + \mathcal{L}_{\text{GAN}}(F, D_X) + \lambda \mathcal{L}_{\text{cyc}}(G, F) ,$$

where λ is a hyperparameter. Although this approach does not exploit the temporal information that is implicit in our demonstration videos, prior work has shown that CycleGAN can nevertheless successfully translate videos frame by frame, such as translating videos of horses into videos of zebras (Zhu et al., 2017).

3.2. Structured Representation Learning

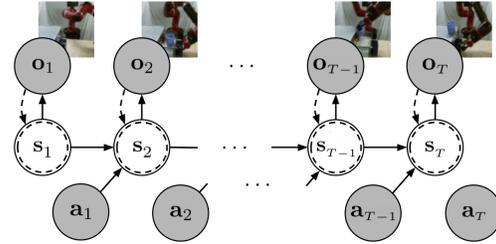


Figure 2. The latent variable model we use to represent robot images and actions, with the generative model depicted in solid lines and the variational family and encoder in dashed lines.

Prior work has shown that, for control learning in image-based domains, state representation learning is an effective tool for improving data efficiency (Hafner et al., 2018; Lee et al., 2019a; Lesort et al., 2018; Zhang et al., 2019). We also use representation learning in our work, and similar to (Lee et al., 2019a), we learn a latent state representation of our image observations by defining a probabilistic temporally-structured latent variable model. In particular, we assume the underlying state of the system s_t is unobserved but evolves as a function of the previous state and action, and we treat the robot images as observations o_t of this state. The generative model for s_t can be summarized as an initial state distribution and a learned neural network dynamics model, i.e.,

$$p(s_1) = \mathcal{N}(s_1; 0, \mathbf{I}) ,$$

$$p(s_{t+1} | s_t, \mathbf{a}_t) = \mathcal{N}(s_{t+1}; \mu(s_t, \mathbf{a}_t), \Sigma(s_t, \mathbf{a}_t)) ,$$

where μ and Σ are parameterized by neural networks. To complete the generative model, we learn a decoder $p(o_t | s_t)$ which we represent as a convolutional neural network. In order to learn this model, we introduce a variational distribution $q(s_{1:T}; o_{1:T})$ which approximates the posterior

$p(\mathbf{s}_{1:T}|\mathbf{o}_{1:T}, \mathbf{a}_{1:T})$, where the $1:T$ notation denotes entire trajectories. We use a mean field variational approximation

$$q(\mathbf{s}_{1:T}; \mathbf{o}_{1:T}) = \prod_t q(\mathbf{s}_t; \mathbf{o}_t),$$

where the encoder $q(\mathbf{s}_t; \mathbf{o}_t)$ is also a convolutional network. We jointly learn the parameters of p and q via maximization of the variational lower bound (ELBO), given by

$$\text{ELBO}[p, q] = \mathbb{E}_q[p(\mathbf{o}_t|\mathbf{s}_t)] - D_{\text{KL}}(q_{\mathbf{s}_1}(\cdot; \mathbf{o}_1)||p_{\mathbf{s}_1}) - \mathbb{E}_q \left[\sum_{t=1}^{T-1} D_{\text{KL}}(q_{\mathbf{s}_{t+1}}(\cdot; \mathbf{o}_{t+1})||p_{\mathbf{s}_{t+1}}(\cdot|\mathbf{s}_t, \mathbf{a}_t)) \right].$$

The entire structured latent variable model is visualized in Figure 2. As we explain in the next section, this model provides us with an encoder and dynamics model that we use for encoding instruction images and for stage-wise model-based planning in the learned latent space.

4. Automated Visual Instruction-Following with Demonstrations

In our problem setting, the task is specified by a set of human demonstration videos, each of which is a trajectory of image observations depicting the human performing the task. We do not assume access to demonstrations given through teleoperation or kinesthetic teaching, which often require specific hardware and human expertise. We also do not assume access to rewards provided through motion capture or other instrumented setups. Thus, our goal is to enable robotic learning while requiring minimal human effort, expertise, and instrumentation. Task specification via a human demonstration is a step toward this goal, and we now detail other design choices that reflect this goal.

4.1. Translation for Goal Concept Acquisition

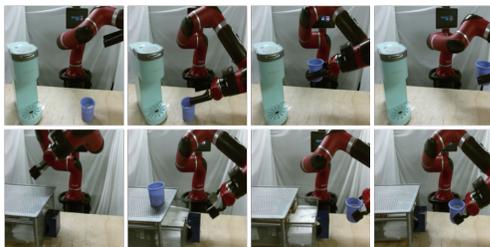


Figure 3. Examples of robot data collected to train the CycleGAN for coffee making (top) and cup retrieval (bottom). Though the robot moves randomly, we cover different settings by, e.g., giving the cup to the robot or opening the drawer.

CycleGAN requires training data from both the source and target domains, and we found that using diverse data was important for generating realistic translations. To that end, the training images we use from the human domain consist

of both human demonstrations as well as a small amount of data in which the human moves around in the scene but does not specifically attempt the task. The training images from the robot domain consist entirely of the robot executing randomly sampled actions in a few different settings, in between which we manually change the environment by shuffling objects or providing the robot with an item (see Figure 3). These human interventions do not require any specialized setups or technical ability, and we found that a handful of interventions was sufficient for CycleGAN data collection (see Section 5). Thus, this overall procedure requires far less human expertise and effort compared to providing a set of robot demonstrations.

Once we have trained the CycleGAN, G provides a mechanism for automatically translating human demonstration videos to robot demonstration videos. Though we could use an IL method that directly learns from the entire translated video, we show in Section 5 that this approach performs poorly. Prior work in IL has shown the efficacy of using select frames or states to imitate rather than using the entire demonstration (Lee et al., 2019b; Pathak et al., 2018; Sharma et al., 2019), and we employ the same approach by extracting and learning from key frames that correspond to the completion of stages. The resulting decoupling of goals from how they can be achieved is especially natural for multi-stage tasks, where achieving particular states that are necessary for overall success is much more important than precisely matching a particular trajectory.

Specifically, the *instruction images* for stage i are taken from the t_i -th frame of the translated videos, where the time steps $\{t_1, \dots, t_S\}$ are manually specified by the user. This is easy for the human as we use a modest number of videos and, as mentioned, the time steps intuitively correspond to the completion of natural stages in the task (see Figure 1). As discussed in Section 2, future work could explore automatic stage discovery for more complex tasks, potentially via video segmentation methods (Jayaraman et al., 2019; Krishnan et al., 2017; Sermanet et al., 2017).

To specify a reward function for each stage, we train success classifiers for each stage $\{\mathcal{C}_1, \dots, \mathcal{C}_S\}$, where \mathcal{C}_i is provided the instruction images at time step t_i as positive examples and all other robot images as negative examples. In order to achieve greater data efficiency, we first encode the images using the encoder $q(\mathbf{s}_t; \mathbf{o}_t)$ and use the learned latent state as input to the classifier. The log probabilities from the classifiers can then be used as a reward signal, similar to (Fu et al., 2018; Singh et al., 2019). Providing rewards in this way does not rely on any manually designed reward shaping or instrumentation. Furthermore, the use of learned classifiers allows for online refinement, an avenue for the human to provide feedback. As we discuss next, this setup lends itself naturally to an automated stage-wise

model-based planning procedure.

4.2. Model-Based RL with Instruction Images

Though we could use any RL algorithm with our classifier-based reward, model-based RL has typically achieved greater data efficiency than model-free methods (Chua et al., 2018; Deisenroth et al., 2014), and the data efficiency afforded by the combination of model-based RL and representation learning further reduces the human supervision required during learning. Thus, the RL procedure that we use is the latent space model-predictive control (MPC) method described in (Zhang et al., 2019), which involves iteratively searching for an optimal sequence of actions using the cross-entropy method (CEM) (Rubinstein & Kroese, 2004), a sampling-based optimization procedure. The search is refined by evaluating model-generated trajectories in the latent space under the reward function, which is given by a classifier in our case. The robot then executes the first action of the optimized sequence. We encapsulate this procedure in the MPC-CEM subroutine in Algorithm 1.

When the robot is attempting stage s , the planner uses the log probability of \mathcal{C}_s as the reward function and aims to surpass a classifier threshold $\alpha \in [0, 1]$, which is a hyperparameter. If this threshold is not met, the planner automatically switches to \mathcal{C}_{s-1} , i.e., it attempts to reset to the beginning of the stage. This forward-reset behavior allows the robot to robustify its performance with very few human-provided resets, as the human only intervenes to fix problems, such as the cup falling over. The robot runs this forward-reset loop for a maximum of K iterations, where K is also a hyperparameter.

Should the threshold be met during planning, the robot will query the human user, at which point the user signals either success or failure through a key press. On failure, the robot switches to the reset behavior, and the loop continues. On success, the robot moves on to the next stage and repeats the same process, with \mathcal{C}_{s+1} specifying the goal and \mathcal{C}_s specifying the reset. This stage-wise learning avoids the compounding errors of trying to learn the entire task all at once. The full procedure is summarized in Algorithm 1.

We include several important improvements on top of this procedure in order to achieve good performance. In line 13 of Algorithm 1, if the human signals failure, the EXPLORE subroutine is invoked, in which the robot moves randomly from the state it ended in and queries the human several times. This provides additional data within the critical area of what \mathcal{C}_s believes is a success. In addition, in line 14, the current image is labeled as a negative for \mathcal{C}_s and added to the dataset along with the images from random exploration. Similar to (Fu et al., 2018; Singh et al., 2019), further training \mathcal{C}_s on this data improves its accuracy by combating the false positive problem and offsetting the artifacts in the translated images that \mathcal{C}_s was initially trained

with. Finally, to further automate the entire process and avoid manual human resets, after the robot completes the last stage, we run Algorithm 1 in reverse in order to get back to the initial state. We repeat this entire process until the robot can reliably complete the entire task.

Algorithm 1 AVID RL Forward Pass

Require: Pre-trained model \mathcal{M} and $\{\mathcal{C}_i\}_{i=0}^S$, initial robot data \mathcal{D}
Require: Max attempts per stage K , classifier threshold α

```

1: stage  $\leftarrow$  1
2: while stage  $\leq S$  do
3:   attempts  $\leftarrow$  0
4:   stage-completed  $\leftarrow$  false
5:   while attempts  $< K$  and not stage-completed do
6:      $\tau \leftarrow$  MPC-CEM( $\mathcal{M}, \mathcal{C}_{\text{stage}}$ );  $\mathcal{D} \leftarrow \mathcal{D} \cup \tau$ 
7:     Increment attempts
8:     if  $\mathcal{C}_{\text{stage}}$ (last observation of  $\tau$ )  $> \alpha$  then
9:       if human signals success then
10:        Add last observation of  $\tau$  to goal images for stage
11:        stage-completed  $\leftarrow$  true
12:       else
13:         $\mathcal{D} \leftarrow \mathcal{D} \cup \text{EXPLORE}()$ 
14:        Train  $\mathcal{M}$  and  $\{\mathcal{C}_i\}_{i=0}^S$  on  $\mathcal{D}$  and goal images
15:       end if
16:       attempts  $\leftarrow$  0
17:       else
18:         $\tau \leftarrow$  MPC-CEM( $\mathcal{M}, \mathcal{C}_{\text{stage}-1}$ );  $\mathcal{D} \leftarrow \mathcal{D} \cup \tau$ 
19:       end if
20:     end while
21:     if stage-completed then
22:       Increment stage
23:     else
24:       Train  $\mathcal{M}$  and  $\{\mathcal{C}_i\}_{i=0}^S$  on  $\mathcal{D}$  and goal images
25:       attempts  $\leftarrow$  0
26:     end if
27:   end while

```

5. Experiments

We aim to answer the following through our experiments:

- (1) Is AVID capable of solving long-horizon, vision-based tasks directly from human demonstrations?
- (2) What benefits, if any, do we gain from using instruction images and latent space planning?
- (3) What costs, if any, do we incur from not having access to demonstrations given directly on the robot?

To answer (1), we evaluate on two complex, long-horizon tasks (see Figure 4). For (2), we compare AVID to an ablation which learns from the entire translated demonstration and one which plans directly in pixel space. Finally, to answer (3), we evaluate IL given ‘‘oracle’’ information, i.e., teleoperated robot demonstrations. We also analyze the human supervision burden of AVID (see Figure 5).

Method	Stage-aware	Demo supervision			Data collected (# images)					
		Human		Robot	Coffee making			Cup retrieval		
		Obs	Obs	Actions	Demo	Pretraining	Online	Demo	Pretraining	Online
AVID (ours)	✓	✓			900	8600	3015	600	7500	1740
Full-video ablation		✓			900	15800	1350	600	14700	1400
Pixel-space ablation	✓	✓			900	15800	0	600	14700	0
TCN (Sermanet et al., 2018)		✓			1395	14400	3060	900	14400	1800
BCO (Torabi et al., 2018)			✓		900	15300	900	900	14400	300
Behavioral Cloning			✓	✓	900	0	0	900	0	0

Table 1. Overview of assumptions and data requirements of the evaluated methods. “Stage-aware” indicates whether the approach allows for leveraging the multi-stage nature of tasks. “Demo supervision” describes the type of demonstration data used: human demonstrations or teleoperated robot demonstrations. We also report the number of images used by each method broken down by demonstrations, pretraining data (including non-demonstration human data and random robot data), and online training data. Because pixel-space models are generally less data efficient than latent variable models, we collect additional pretraining data for all methods except for ours. Note that, analogous to prior work (Ebert et al., 2018), the pixel-space ablation is not trained online because it is too computationally expensive, and behavioral cloning does not have an online training phase.

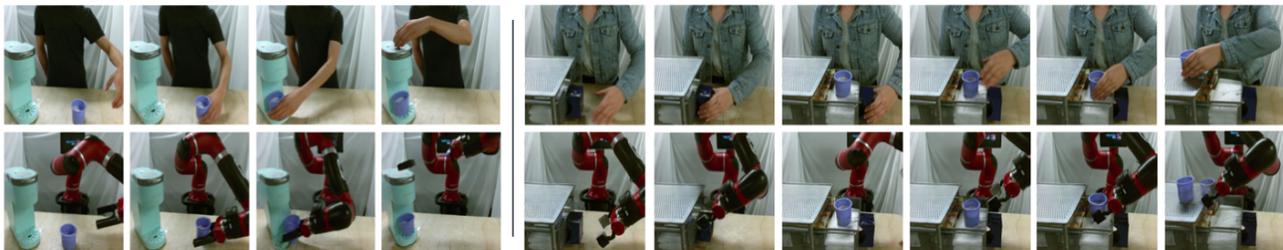


Figure 4. Sample sequence of instructions for each of the two tasks, coffee making (left) and cup retrieval (right). The instruction images are segmented from a human demonstration (top), then translated into the robot’s domain (bottom) via CycleGAN. Note the artifacts in the generated translations, e.g., the displaced robot gripper in the last instruction image on the left and the deformed cup in the last image on the right.

5.1. Comparisons

Behavioral cloning. Behavioral cloning directly applies supervised learning on observation-action pairs from an expert to approximate the optimal policy. Although this “oracle” approach uses teleoperated demonstrations given on the robot, it is not designed to incorporate the notion of stages and is known to suffer from distributional shift arising from compounding errors (Ross et al., 2011).

Behavioral Cloning from Observation (BCO) (Torabi et al., 2018). BCO also uses teleoperated demonstrations but only assumes access to observations. This method infers the actions from the observations using a learned inverse model, trained separately, and then performs behavioral cloning on the inferred actions.

Full-video ablation. We consider an ablation that uses entire demonstrations instead of instruction images. A natural way to do so is to use an IL method that does not require expert actions (Edwards et al., 2019; Liu et al., 2018; Pathak et al., 2018; Sun et al., 2019; Torabi et al., 2018; Yu et al., 2018). We use BCO, using translated human demonstrations rather than teleoperated demonstrations.

Pixel-space ablation. Deep visual foresight (DVF) (Ebert

et al., 2018) is a visual model-based RL method that performs planning in pixel space instead of a learned latent space. We consider an ablation of AVID in which we use translated instruction images and our stage-wise training procedure but replace latent space planning with DVF.

Time-contrastive networks (Sermanet et al., 2018). We evaluate the single-view version of TCN, which learns an embedding via a temporal consistency loss in order to generalize from human demonstrations to robot execution. For fairness and ease of comparison, we use the same RL subroutine as AVID. While AVID learns classifiers for the reward function, this comparison uses the negative Euclidean distance between TCN embeddings of the human demonstration and model-generated trajectories, decoded using $p(\mathbf{o}_t | \mathbf{s}_t)$, as the reward function. We compare to the single-view version of TCN, as the multi-view version requires filming from multiple viewpoints, and we do not make the same assumption in our setup.

5.2. Experimental Setup

We evaluate our method on two temporally-extended tasks from vision: operating a personal coffee machine and retrieving a cup from a closed drawer. These tasks illustrate

Supervision	Method	Coffee making			Cup retrieval				
		Stage 1	Stage 2	Stage 3	Stage 1	Stage 2	Stage 3	Stage 4	Stage 5
Human Demos	AVID (ours)	100%	80%	80%	100%	100%	100%	80%	70%
	Full-video ablation	70%	10%	0%	0%	0%	0%	0%	0%
	Pixel-space ablation	60%	20%	0%	50%	50%	30%	10%	0%
	TCN (Sermanet et al., 2018)	10%	10%	0%	60%	20%	0%	0%	0%
Teleoperated Robot Demos	BCO (Torabi et al., 2018)	80%	30%	0%	30%	10%	0%	0%	0%
	Behavioral Cloning	90%	90%	90%	100%	100%	60%	60%	40%

Table 2. We report success rates up to and including each stage for both tasks, over 10 trials. The top rows are methods that learn from human demonstrations, and we bold the best performance in this category. The bottom two are methods trained with demonstrations obtained via teleoperating the robot. AVID outperforms all other methods from human demonstrations, succeeding 8/10 times on coffee making and 7/10 times on cup retrieval, and even outperforms behavioral cloning on the later stages of cup retrieval.

that our method can learn to sequentially compose several skills by following a set of instructions extracted from a human demonstration. For all our experiments, we use end-effector velocity control on a Sawyer robotic manipulator, and our observations consist only of 64x64x3 RGB images.

For data collection, we film videos of human actions and random robot trajectories. This gives us data from the two domains, and we train a CycleGAN using this data for AVID, the full-video ablation, and pixel-space ablation. The human demonstration videos are then reused to generate translated robot demonstrations. We also reuse the random robot trajectories to train the models needed for each method: the latent variable model for AVID, the dynamics model for the pixel-space ablation, the inverse models for the full-video ablation and BCO, and the embedding function for TCN. A full breakdown can be found in Table 1.

Coffee making. The coffee making task has three stages as depicted in Figure 4. Starting with the cup on the table, the instructions are to pick up the cup, place it in the machine, and press the button on the top. For this task, we used 30 human demonstrations (900 images) along with 500 images of random human data, which took about 20 minutes. We also collected trajectories of the robot executing randomly sampled actions, with the human changing the setting 6 times by moving the cup and placing the cup in the robot’s gripper. We set the classifier threshold α to 0.8 and the maximum number of attempts per stage K to 3.

Cup retrieval. Cup retrieval has five stages, as shown in Figure 4: grasping the drawer handle, opening the drawer, moving the arm out of the way, picking up the cup, and placing it on top of the drawer. We found that the intermediate stage of moving the arm was important for good performance, as the model-based planning would otherwise bump into the drawer door. Providing this instruction was trivial with our setup, as we simply specified an additional instruction time step between the “natural” instructions. For this task, we used 20 human demonstrations (600 images) and 300 images of random human data, again amounting to about 20 minutes of human time. In collect-

ing the random robot data, the human changed the setting 11 times by moving the cup, placing the cup in the robot’s gripper, and opening the drawer. α and K were again set to 0.8 and 3.

5.3. Experimental Results

We report success rates for both coffee making and cup retrieval in Table 2. Success rates are evaluated using 10 trials per method per task, with no human feedback or manual intervention. Success metrics are defined consistently across all methods and are strict, e.g., placing the cup on top of the drawer but toppling the cup over is a failure. The supplementary video shows representative evaluation trials, labeled with success and failure, for each method.¹ For the final stage of coffee making, as it is difficult to visually discern whether the robot has noticeably depressed the coffee machine button, success is instead defined as being within 5 cm above the button, such that a pre-programmed downward motion always successfully presses the button.

For both tasks, AVID achieves the best performance among the methods that use human demonstrations, consistently learning the earlier stages perfectly and suffering less from accumulating errors in the later stages. As seen in the supplementary video, the failures of AVID for both tasks correspond to generally good behavior with small, but significant, errors, such as knocking over or narrowly missing the cup. AVID makes constant use of automated resetting and retrying during both RL and the final evaluation, allowing for more robust behavior compared to the methods that cannot exploit the stage-wise decomposition of the tasks. Though the pixel-space ablation also incorporates resetting, this method is less successful than AVID. This may be due to the fact that the pixel-space model is not trained online or used for replanning as this is prohibitively expensive. This highlights the benefits of our latent variable model and planning procedure.

The full-video ablation can learn to pick up the cup 70% of

¹<https://sites.google.com/view/rss20avid>

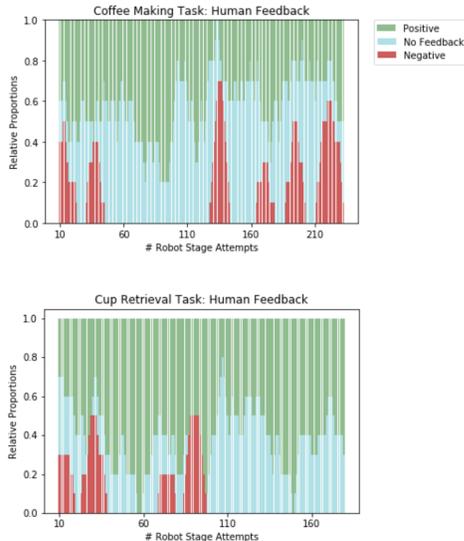


Figure 5. Visualizing human feedback during the learning process for coffee making (top) and cup retrieval (bottom). The x-axis is the total number of robot stage attempts, and the y-axis indicates the proportions of feedback types, smoothed over the ten most recent attempts. “No feedback” means that the classifier did not signal success and the robot automatically switched to resetting. Coffee making and cup retrieval used a total of 131 and 126 human key presses, respectively.

the time for the coffee making task, but it fails on the later stages and also fails entirely for cup retrieval. This indicates the difficulty of learning from full translated demonstrations due to the artifacts in the translations and the absence of stage-wise training. The supplementary video shows that, although the behavior of this method visually looks almost correct in many instances, it generally still fails in subtle ways. Finally, despite using publicly available code from the authors, we were unable to achieve good performance with TCN on coffee making, though TCN was moderately successful on the early stages of cup retrieval. We note that we used the single-view version of this method, and the authors also found that this version did not perform well for their tasks (Sermanet et al., 2018). As previously mentioned, the multi-view version of TCN is not applicable in our case as we do not assume multiple views of human demonstrations at training time.

In order to understand how using translated human demonstrations compares to using real robot demonstrations, we evaluate BCO and behavioral cloning from demonstrations obtained through teleoperation. As expected, BCO performs better in this case than from translated demonstrations, however the performance is still significantly worse than AVID from translated demonstrations. This indicates that learning from full demonstrations, even directly obtained from the robot, suffers from accumulating errors for

long-horizon tasks. We note that, to our knowledge, BCO has never been tested on multi-stage tasks from image observations (Torabi et al., 2018).

Finally, behavioral cloning outperforms our method for coffee making but surprisingly performs worse on cup retrieval. We believe that this is because the cup retrieval task has more stages and is harder to learn without explicit stage-wise training. Thus, compared to AVID, behavioral cloning has more stringent assumptions and may perform worse for more complex tasks, but the drawback of AVID is that it uses an RL training phase that requires human feedback. To understand the human burden associated with this training procedure, we plot the human feedback we provide during training in Figure 5. Training took about one hour for each task, and in both cases AVID used less than 150 key presses to learn the task, which was easy for a human supervisor to provide. We view this as a practical way to enable robots to learn complex tasks.

6. Discussion and Future Work

We presented AVID, a method for learning visual robotic multi-stage tasks directly from videos of human demonstrations. AVID uses image-to-image translation, via CycleGAN, to generate robot instruction images that enable a stage-wise model-based RL algorithm to learn from raw image observations. The RL algorithm is based on latent space model-based planning and incorporates learned resets and human feedback to learn temporally extended tasks robustly and efficiently. We demonstrated that AVID is capable of learning multi-stage coffee making and cup retrieval tasks, attains substantially better results than prior methods and ablations using human demonstrations, and even outperforms behavioral cloning using teleoperated demonstrations on one of the tasks.

The most exciting direction for future work is to amortize the cost of data collection and CycleGAN training across multiple tasks, rather than training separate CycleGANs for each task of interest. For example, the cup retrieval CycleGAN can also be used to translate demonstrations for placing the cup back in the drawer, meaning that this new task can be learned from a handful of human demonstrations with no additional upfront cost. We can generalize this by training a CycleGAN on an initial large dataset, e.g., many different human and robot behaviors in a kitchen that has a coffee machine, multiple drawers, and numerous other objects. This should enable any new task in the kitchen to be learned with just a few human demonstrations of the task, and this is a promising direction toward truly allowing robots to learn by watching humans.

References

- Akgun, B., Cakmak, M., Yoo, J., and Thomaz, A. Trajectories and keyframes for kinesthetic teaching: A human-robot interaction perspective. In *HRI*, 2012.
- Argall, B., Chernova, S., Veloso, M., and Browning, B. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 2009.
- Billard, A. and Mataric, M. Learning human arm movements by imitation: Evaluation of a biologically inspired connectionist architecture. *Robotics and Autonomous Systems*, 2001.
- Calinon, S., Evrard, P., Gribovskaya, E., Billard, A., and Kheddar, A. Learning collaborative manipulation tasks by demonstration using a haptic interface. In *ICAR*, 2009.
- Christiano, P., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences. In *NIPS*, 2017.
- Chua, K., Calandra, R., McAllister, R., and Levine, S. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *NIPS*, 2018.
- Daniel, C., Viering, M., Metz, J., Kroemer, O., and Peters, J. Active reward learning. In *RSS*, 2014.
- Deisenroth, M., Fox, D., and Rasmussen, C. Gaussian processes for data-efficient learning in robotics and control. *PAMI*, 2014.
- Ebert, F., Finn, C., Dasari, S., Xie, A., Lee, A., and Levine, S. Visual foresight: Model-based deep reinforcement learning for vision-based robotic control. *arXiv preprint arXiv:1812.00568*, 2018.
- Edwards, A., Sahni, H., Schroecker, Y., and Isbell, C. Imitating latent policies from observation. In *ICML*, 2019.
- Fu, J., Singh, A., Ghosh, D., Yang, L., and Levine, S. Variational inverse control with events: A general framework for data-driven reward definition. In *NIPS*, 2018.
- Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., and Davidson, J. Learning latent dynamics for planning from pixels. In *ICML*, 2018.
- Han, W., Levine, S., and Abbeel, P. Learning compound multi-step controllers under unknown dynamics. In *IROS*, 2015.
- Hayes, G. and Demiris, J. A robot controller using learning by imitation. 1994.
- Jayaraman, D., Ebert, F., Efros, A., and Levine, S. Time-agnostic prediction: Predicting predictable video frames. In *ICLR*, 2019.
- Kalakrishnan, M., Pastor, P., Righetti, L., and Schaal, S. Learning objective functions for manipulation. In *ICRA*, 2013.
- Kalashnikov, D., Irpan, A., Pastor, P., Ibarz, J., Herzog, A., Jang, E., Quillen, D., Holly, E., Kalakrishnan, M., Vanhoucke, V., and Levine, S. QT-Opt: Scalable deep reinforcement learning for vision-based robotic manipulation. In *CoRL*, 2018.
- Kober, J. and Peters, J. Learning motor primitives for robotics. In *ICRA*, 2009.
- Kober, J. and Peters, J. Policy search for motor primitives in robotics. *Machine Learning*, 2011.
- Kober, J., Bagnell, J., and Peters, J. Reinforcement learning in robotics: A survey. *IJRR*, 2013.
- Kohl, N. and Stone, P. Policy gradient reinforcement learning for fast quadrupedal locomotion. In *ICRA*, 2004.
- Kormushev, P., Calinon, S., and Caldwell, D. Robot motor skill coordination with EM-based reinforcement learning. In *IROS*, 2010.
- Krishnan, S., Garg, A., Patil, S., Lea, C., Hager, G., Abbeel, P., and Goldberg, K. Transition state clustering: Unsupervised surgical trajectory segmentation for robot learning. *IJRR*, 2017.
- Kuniyoshi, Y., Inaba, M., and Inoue, H. Learning by watching: Extracting reusable task knowledge from visual observation of human performance. *IEEE Transactions on Robotics and Automation*, 1994.
- Kurutach, T., Tamar, A., Yang, G., Russell, S., and Abbeel, P. Learning plannable representations with causal infoGAN. *arXiv preprint arXiv:1807.09341*, 2018.
- Lee, A., Nagabandi, A., Abbeel, P., and Levine, S. Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model. *arXiv preprint arXiv:1907.00953*, 2019a.
- Lee, K., Su, Y., Kim, T., and Demiris, Y. A syntactic approach to robot imitation learning using probabilistic activity grammars. *Robotics and Autonomous Systems*, 2013.
- Lee, Y., Hu, E., Yang, Z., and Lim, J. To follow or not to follow: Selective imitation learning from observations. *arXiv preprint arXiv:1912.07670*, 2019b.
- Lesort, T., Díaz-Rodríguez, N., Goudou, J., and Filliat, D. State representation learning for control: An overview. *Neural Networks*, 2018.

- Liu, M., Breuel, T., and Kautz, J. Unsupervised image-to-image translation networks. In *NIPS*, 2017.
- Liu, Y., Gupta, A., Abbeel, P., and Levine, S. Imitation from observation: Learning to imitate behaviors from raw video via context translation. In *ICRA*, 2018.
- Manschitz, S., Kober, J., Gienger, M., and Peters, J. Learning to sequence movement primitives from demonstrations. In *IROS*, 2014.
- Miyamoto, H., Schaal, S., Gandolfo, F., Gomi, H., Koike, Y., Osu, R., Nakano, E., Wada, Y., and Kawato, M. A Kendama learning robot based on bi-directional theory. *Neural Networks*, 1996.
- Nair, A., Chen, D., Agrawal, P., Isola, P., Abbeel, P., Malik, J., and Levine, S. Combining self-supervised learning and imitation for vision-based rope manipulation. In *ICRA*, 2017.
- Nehaniv, C. and Dautenhahn, K. The correspondence problem. *Imitation in Animals and Artifacts*, 2002.
- OpenAI et al. Learning dexterous in-hand manipulation. *arXiv preprint arXiv:1808.00177*, 2018.
- Pastor, P., Righetti, L., Kalakrishnan, M., and Schaal, S. Online movement adaptation based on previous sensor experiences. In *IROS*, 2011.
- Pathak, D., Mahmoudieh, P., Luo, G., Agrawal, P., Chen, D., Shentu, Y., Shelhamer, E., Malik, J., Efros, A., and Darrell, T. Zero-shot visual imitation. In *ICLR*, 2018.
- Ramirez-Amaro, K., Beetz, M., and Cheng, G. Transferring skills to humanoid robots by extracting semantic representations from observations of human activities. *Artificial Intelligence*, 2017.
- Rhinehart, N. and Kitani, K. First-person activity forecasting with online inverse reinforcement learning. In *ICCV*, 2017.
- Ross, S., Gordon, G., and Bagnell, J. A reduction of imitation learning and structured prediction to no-regret online learning. In *AISTATS*, 2011.
- Rubinstein, R. and Kroese, D. *The Cross Entropy Method: A Unified Approach To Combinatorial Optimization, Monte-Carlo Simulation (Information Science and Statistics)*. Springer-Verlag New York, Inc., 2004.
- Rusu, A., Večerík, M., Rothörl, T., Heess, N., Pascanu, R., and Hadsell, R. Sim-to-real robot learning from pixels with progressive nets. In *CoRL*, 2017.
- Schaal, S., Ijspeert, A., and Billard, A. Computational approaches to motor learning by imitation. *Philosophical Transaction of the Royal Society of London, Series B*, 2003.
- Schenck, C. and Fox, D. Visual closed-loop control for pouring liquids. In *ICRA*, 2017.
- Sermanet, P., Xu, K., and Levine, S. Unsupervised perceptual rewards for imitation learning. In *RSS*, 2017.
- Sermanet, P., Lynch, C., Chebotar, Y., Hsu, J., Jang, E., Schaal, S., and Levine, S. Time-contrastive networks: Self-supervised learning from video. In *ICRA*, 2018.
- Sharma, P., Pathak, D., and Gupta, A. Third-person visual imitation learning via decoupled hierarchical control. In *NeurIPS*, 2019.
- Singh, A., Yang, L., Hartikainen, K., Finn, C., and Levine, S. End-to-end robotic reinforcement learning without reward engineering. In *RSS*, 2019.
- Stulp, F., Theodorou, E., and Schaal, S. Reinforcement learning with sequences of motion primitives for robust manipulation. *IEEE Transactions on Robotics*, 2012.
- Sun, W., Vemula, A., Boots, B., and Bagnell, J. Provably efficient imitation learning from observation alone. In *ICML*, 2019.
- Torabi, F., Warnell, G., and Stone, P. Behavioral cloning from observation. In *IJCAI*, 2018.
- Tow, A., Sünderhauf, N., Shirazi, S., Milford, M., and Leitner, J. What would you do? Acting by learning to predict. In *IROS*, 2017.
- van Hoof, H., Hermans, T., Neumann, G., and Peters, J. Learning robot in-hand manipulation with tactile features. In *Humanoids*, 2015.
- Yahya, A., Li, A., Kalakrishnan, M., Chebotar, Y., and Levine, S. Collective robot reinforcement learning with distributed asynchronous guided policy search. In *IROS*, 2017.
- Yang, Y., Li, Y., Fermuller, C., and Aloimonos, Y. Robot learning manipulation action plans by “watching” unconstrained videos from the world wide web. In *AAAI*, 2015.
- Yu, T., Finn, C., Xie, A., Dasari, S., Abbeel, P., and Levine, S. One-shot imitation from observing humans via domain-adaptive meta-learning. In *RSS*, 2018.
- Zhang, M., Vikram, S., Smith, L., Abbeel, P., Johnson, M., and Levine, S. SOLAR: Deep structured representations for model-based reinforcement learning. In *ICML*, 2019.

Zhang, T., McCarthy, Z., Jow, O., Lee, D., Goldberg, K., and Abbeel, P. Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. In *ICRA*, 2018.

Zhu, J., Park, T., Isola, P., and Efros, A. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.