

Manipulate by Seeing: Creating Manipulation Controllers from Pre-Trained Representations

Jianren Wang* Sudeep Dasari* Mohan Kumar Srirama Shubham Tulsiani Abhinav Gupta

Carnegie Mellon University

Abstract

The field of visual representation learning has seen explosive growth in the past years, but its benefits in robotics have been surprisingly limited so far. Prior work uses generic visual representations as a basis to learn (task-specific) robot action policies (e.g., via behavior cloning). While the visual representations do accelerate learning, they are primarily used to encode visual observations. Thus, action information has to be derived purely from robot data, which is expensive to collect! In this work, we present a scalable alternative where the visual representations can help directly infer robot actions. We observe that vision encoders express relationships between image observations as distances (e.g., via embedding dot product) that could be used to efficiently plan robot behavior. We operationalize this insight and develop a simple algorithm for acquiring a distance function and dynamics predictor, by fine-tuning a pre-trained representation on human collected video sequences. The final method is able to substantially outperform traditional robot learning baselines (e.g., 70% success v.s. 50% for behavior cloning on pick-place) on a suite of diverse real-world manipulation tasks. It can also generalize to novel objects, without using any robot demonstrations during train time. For visualizations of the learned policies please check: <https://agi-labs.github.io/manipulate-by-seeing/>.

1. Introduction

The lack of suitable, large-scale data-sets is a major bottleneck in robot learning. Due to the physical nature of data collection, robotics data-sets are: (a) hard to scale; (b) collected in sterile, non-realistic environments (e.g. robotics lab); (c) too homogeneous (e.g. toy objects with fixed backgrounds/lighting). In contrast, vision data-sets contain diverse tasks, objects,

*Denotes equal contribution.

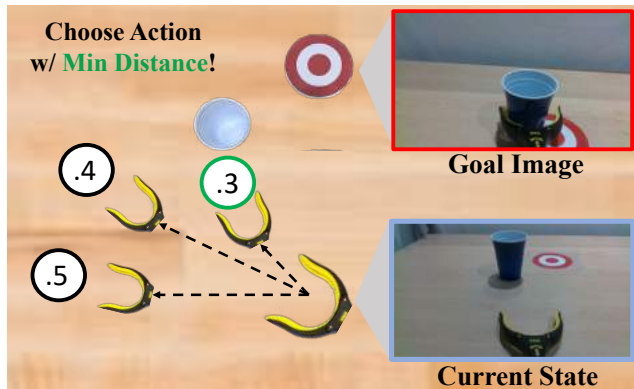


Figure 1: This paper proposes to solve a range of manipulation tasks (e.g. pushing) by learning a functional distance metric within the embedding space of a pre-trained network. This distance function – in combination with a learned dynamics model – can be used to greedily plan for robot actions that reach a goal state. Our experiments reveal that the proposed method can outperform SOTA robot learning methods across four diverse manipulation tasks.

and settings (e.g. Ego4D [14]). Therefore, recent approaches have explored transferring priors from large scale vision data-sets to robotics settings. What is the right way to accomplish this?

Prior work uses vision data-sets to pre-train representations [31, 29, 32] that encode image observations as state vectors (i.e. $s = R(i)$). This visual representation is then simply used as an input for a controller learned from robot data – e.g. a policy $\pi(a|s)$, trained with expert data via Behavior Cloning [35], or a Value function $V(s)$, trained using exploratory roll-outs via Reinforcement Learning [43]. Is this approach the most efficient way to use pre-trained representations? We argue that pre-trained networks can do more than just represent states, since their latent space already encodes semantic, task-level information – e.g. by plac-

ing semantically similar states more closely together. Leveraging this structure to infer actions, could enable us to use significantly less robotic data during train time.

Our paper achieves this by fine-tuning a pre-trained representation into: (a) a one-step dynamics module, $F(s, a)$, that predicts how the robot’s next state given the current state/action; and (b) a “functional distance module”, $d(s, g)$, that calculates how close the robot is to achieving its goal g in the state s . The distance function is learned with limited human demonstration data, using a contrastive learning objective (see Fig. 2). Both d and F are used in conjunction to greedily plan robot actions (see Fig. 1 for intuition). Our experiments demonstrate that this approach works better than policy learning (via Behavior Cloning), because the pre-trained representation itself does the heavy lifting (thanks to its structure) and we entirely dodge the challenge of multi-modal, sequential action prediction. Furthermore, our learned distance function is both stable and easy to train, which allows it to easily scale and generalize to new scenarios.

To summarize, we show a simple approach for exploiting the information hidden in pre-trained visual representations. Our contributions include:

- Developing a simple algorithm for acquiring a distance function and dynamics model by fine-tuning a pre-trained visual representation on minimal (human collected) data.
- Creating an effective manipulation controller that substantially outperforms State-of-the-Art (SOTA) prior methods from the robot learning community (e.g. Behavior Cloning [35], Offline-RL [20], etc.).
- Demonstrating that our approach can handle four realistic manipulation tasks, generalize to new objects and settings, and solve challenging scenarios w/ multi-modal action distributions.

2. Related Work

Behavior Cloning This paper adopts the Learning from Demonstration (LfD) problem setting [4, 7, 36], where a robot must acquire manipulation behaviors given expert demonstration trajectories. A standard approach in this space is to learn a policy (π) via Behavior Cloning [35] (BC), which directly optimizes π to match the expert’s action distribution. While conceptually simple, realistic action distributions are difficult to model, since they are inherently multi-modal and even small errors compound over time. Thus, policy

learning requires extensive intensive network engineering (e.g. transformer architectures [10, 9], multi-modal prediction heads [41, 27], etc.) and/or human-in-the-loop data collection algorithms [35, 18] to work in practice. Instead of learning a policy, we learn a *functional distance metric* that captures the how “close” a state is to reaching a target goal. This lets us build a manipulation controller using a simple greedy planner during test time, without any explicit action prediction!

Offline RL Broadly speaking, the field of Reinforcement Learning [43] (RL) seeks to learn a value/advantage/Q function by assuming access to a reward signal, which “evaluates” a state (e.g. +1 reward for reaching goal). RL algorithms usually require environment interaction to learn, though the field of Offline-RL [25] seeks to extend these systems to learn from offline trajectories. While these approaches have created some impressive robotics demos ranging from locomotion [24, 22] to dexterous manipulation [2, 16, 11], RL methods are data hungry [33, 2], difficult to implement, require extreme “reward engineering” (e.g. Meta-World [49] reward functions are 50 lines!), hampered by unstable learning dynamics [21], and poorly generalize in realistic robotic manipulation settings [12]. In contrast, our method learns a proxy for value functions (i.e. distances) purely from offline data, using representation learning algorithms (from the vision field) that are much more stable.

Learning Visual Rewards Our distance learning approach is analogous to learning visual rewards. Prior work learned success classifiers [40, 47, 46] and/or video similarity metrics [39, 5, 37, 6, 8] from expert video demonstrations. Once learned, these modules were used to parameterize reward functions that could be used to train policies [1], and/or as cost functions for planners [3]. However, these papers mostly consider simple settings (e.g. simulation) or require the test setting to exactly match train time. This is because the learned reward functions are often noisy and/or poorly generalize to new scenes. In contrast, our learned distance metric is stable and can easily adapt to new scenarios, thanks to the pre-trained network’s latent structure. This allows us to solve diverse tasks during test time on a real robot, using a simple (and fast) shooting method planner.

Distance Learning for Navigation Finally, prior work from Hahn et. al. [15] explored distance learning to solve visual navigation tasks from offline video data. Our work is inspired from this approach, but makes

three key innovations in order to apply it to (more challenging) robotic manipulation tasks. First, we simplify the network architecture to operate in representation space instead of using a graph network. Second, we learn the distances using a self-supervised contrastive loss, instead of relying on hand-defined objectives. And third, we add a dynamics function that proves to be critical for effective manipulation (see Table. 3).

3. Methods

Preliminaries This paper considers learning goal-conditioned manipulation behaviors from image observations. The robot agent is provided a goal observation (I_g) – e.g. target object in robot gripper – and an observation (I_t) for the current time-step t . The robot must process these observations and decide an action to take (a_t). Note that all observations (I) are wrist-mounted RGB camera images with no depth or proprioceptive data. Additionally, the actions (a) are specified as arbitrary $SE(3)$ transforms for the robot’s end-effector (more in Supplement B.1). Our goal is to learn a robotic manipulation controller, using a set of training trajectories $\mathcal{D} = \{\tau_1, \dots, \tau_N\}$, where $\tau_i = \{I_g, I_1, a_1, \dots, a_{T-1}, I_T\}$. The test and train settings are visualized in Fig. 4.

Our Approach Our method leverages a pre-trained representation network, R , to encode observations, $i_t = R(I_t)$, and enable control via distance learning. Specifically, we use contrastive representation learning methods [30, 28] to learn a distance metric, $d(i_j, i_k)$, within the pre-trained embedding space. The key idea is to use this distance metric to select which of the possible future state is closest to the goal state. But how do we predict possible future states? We explicitly learn a dynamics function, $F(i_t, a_t)$ that predicts future state for a possible action a_t . During test time, we predict multiple future states using different possible action and select the one which is closes to goal state. The following sections describe the learned dynamics module (see Sec. 3.1), distance learning method (see Sec. 3.2), and test-time controller for robot deployment (see Sec. 3.4) in detail. Fig. 2 provides a visual depiction of our training algorithm. Pseudo-code and hyper-parameters are presented in Supplement A.3.

3.1. Dynamics Prediction

An ideal dynamics function would perfectly capture how a robot’s actions effects its environment. In our setting, this translates to predicting the next observation embedding, given the current embedding and commanded action: $F(i_t, a_t) = i_{t+1}$. Thus, F can

be learned via regression by minimizing reconstruction loss in embedding space: $\mathcal{L}_F = \|F(i_t, a_t) - i_{t+1}\|_2$ (see Fig. 2, left). This module has two primary uses: (1) during training it acts as a physically grounded regularizer that relays action information into the embedding space, and (2) during test time it enables the robot to plan actions in embedding space. As you will see, both properties are leveraged extensively in the rest of our method.

3.2. Learning Task-Centric Distances

Our distance module $d(i_j, i_k)$ seeks to learn functional distances that encode task-centric reasoning (e.g. must reach for object before pushing it) alongside important physical priors (e.g. object is to the left so move left). How can we learn such a distance space? Since we have access to expert trajectories τ , we know that the next state sequentially visited by the expert (i_{t+1}) is closer to the goal state (i_g) than arbitrary states (\hat{i}) reachable from i_t – i.e. $d(i_{t+1}, i_g) \ll d(\hat{i}, i_g)$.

Our insight is that a distance metric with these properties can be learned via contrastive learning. Specifically, we define $d(i_j, i_k) = -\cos(i_j, i_k)$, where \cos is cosine similarity, and sample a observation-state-goal tuples (i_t, a_t, i_g), alongside random noise action candidates $\hat{a}^1, \dots, \hat{a}^n$ sampled from \mathcal{D} . We apply NCE loss and get:

$$\mathcal{L}_d = \frac{\exp(-d(F(i_t, a_t), i_g))}{\exp(-d(F(i_t, a_t), i_g)) + \sum_j \exp(-d(F(i_t, \hat{a}^j), i_g))}$$

This loss creates a warped embedding space where i_{t+1} (hallucinated by $F(i_t, a_t)$) is pushed closer towards the goal state i_g , than other arbitrary states reachable from i_t (again hallucinated by F). This process is shown visually in Fig. 2 (right), and precisely satisfies our criteria for good distance functions.

3.3. Training Details

Both modules are trained jointly resulting in a final loss: $\mathcal{L} = \lambda_d \mathcal{L}_d + \lambda_F \mathcal{L}_F$. Note that F is implemented as small neural networks with 2 hidden layer, while d does not require any extra learned networks since its implemented via contrastive loss in the metric space. The shared representation network, R , we use is ResNet-18 initialized by R3M weights. The ADAM [19] stochastic gradient descent optimizer and back-propagation are used to train the network end-to-end. Please refer to Supplement A.3. for additional details.

3.4. Test Time Robot Deployment

During test time our learned modules must be able to solve real manipulation tasks when deployed on

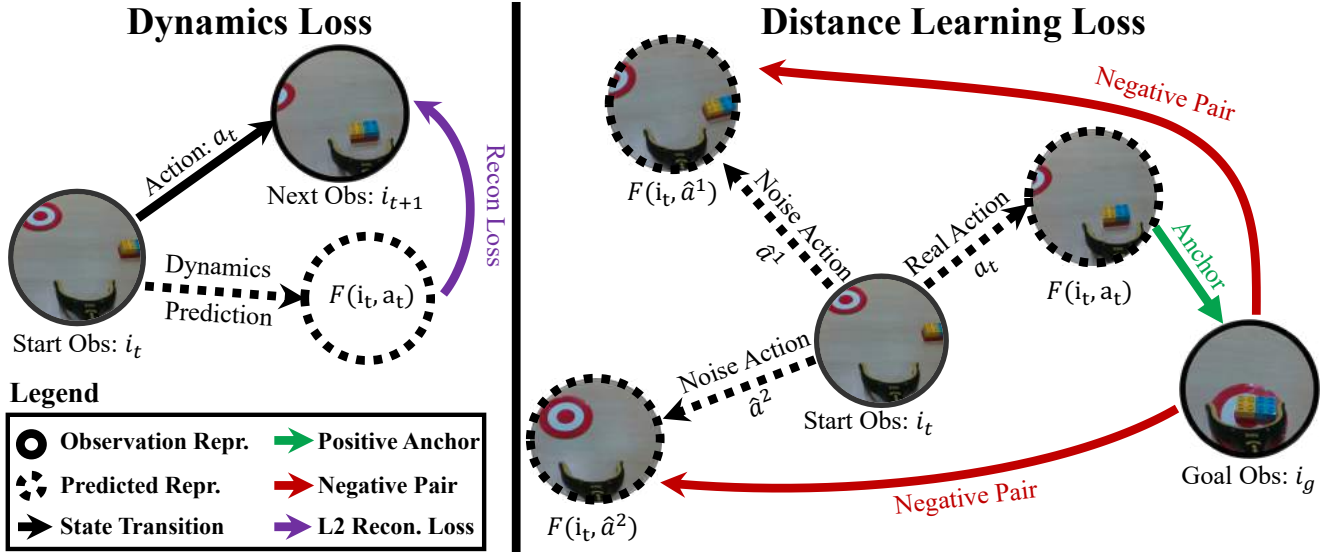


Figure 2: We visualize the loss functions used to train our method. The dynamics function is trained via reconstruction loss in embedding space (left). The distance function is trained via contrastive learning, with positive anchors chosen by predicting the next state using the ground truth action, $F(i_t, a_t)$, and predicting negative pairs chosen using noisy actions $F(i_t, \hat{a}^j)$.

robot hardware. Thus, we develop a simple inference policy that solves for robot actions using our learned distance and dynamics function. First, a cost function $C_g(i) = d(i, i_g)$ is parameterized given a goal observation I_g and the learned distance function. C_g encodes how far an arbitrary image observation is from the goal image. The optimal action, a_t^* , will take the robot closest to the goal, hence: $a_t^* = \mathop{\text{argmin}}_a C_g(F(i_t, a))$. Note that we use the dynamics function F to predict the next state, since we do not know the real i_{t+1}^* during test time. The cost minimization is performed using the shooting method: random candidate actions, a_1, \dots, a_N , are sampled from \mathcal{D} , passed through the dynamics function, and the action with minimum cost $C_g(F(i_t, a_i))$ is executed on the robot. This process is illustrated in Fig. 1. Policy execution ends once the distance between the current state and the goal falls below a threshold value: $C_g(i_t) < \lambda$. While a more complicated planner could’ve been used, this solution was adopted due to its speed, simplicity, lack of hyperparameters, and good empirical performance.

The final detail to discuss is gripper control for prehensile tasks (e.g. pick and place). Since the gripper status is binary (open/close) and highly correlated with the current state (i.e. close when object in hand), it makes more sense to handle it implicitly rather than as part of the action command a_t . Thus, we trained a gripper action classifier $\mathcal{G}(i_t) \in [0, 1]$ that predicts the probability of closing the gripper given the current

image embedding. This is also implemented by adding a single layer to ResNet-18 initialized by R3M weights and can be trained using a small amount (< 100) of human labelled images from the train dataset \mathcal{D} . As a result, our system can seamlessly handle gripper control in prehensile tasks, without requiring gripper labels for every frame in \mathcal{D} !

4. Experimental Setup

Our method is tested on four different manipulation tasks (see Fig. 3) that require a mix of *high-level reasoning* (e.g. go to object before target), *low level precision* (must carefully grab knob to turn it), and *time-time generalization* (e.g. push novel object). We collect behavior data for each task from human demonstrators as described below. The trained policies are deployed on a real Franka Pandas robot (see Fig. 4, right). Additional details on the hardware setup and control stack are presented in Supplement B.1.

4.1. Tasks

We now introduce the tasks used in our experiments (pictured in Fig. 3), and provide more details on their train and test conditions:

Pushing For this task (see Fig. 3a) the robot must approach a novel object placed on the table and push it to the goal location (marked by target printout).

Training: Our method was trained on a dataset of



(a) Pushing (b) Pick and Place (c) Door Opening (d) Knob Turning

Figure 3: Our method is tested on 4 different manipulation tasks (pictured above). These tasks test different skill axis, ranging from task-level reasoning (e.g. detect target, move to object before goal, etc.) to fine-grained motor control (e.g. grab top of knob to turn).

Low-Cost Human Videos Transferred to Robot




Figure 4: In our problem setting we use a low-cost reacher grabber tool (left) to collect training demonstrations. These demonstrations are used to acquire a robot controller purely through distance/representation learning. The final system is deployed on a robot (right) to solve various tasks at test-time.

100 demonstrations with diverse objects and randomized targets. **Testing:** During test time, the method was evaluated using 20 trails that involved unseen objects and new target printouts placed in randomly sampled positions. A trail is deemed successful if the robot pushes the object onto the target. This is the only task that did not require gripper control.

Pick and Place This task (see Fig. 3b) is analogous to the pushing task (described above), except the robot must now pick the object up from its initial location and place it in a target bowl. This task requires more precision than pushing, since grasping an object is harder than localizing it. **Training:** We collect a dataset of 400 demonstrations for training, using randomized train objects and target bowls. **Testing:** The method is evaluated using 20 trails that used novel objects and unseen target bowls (i.e. analogous to pushing). A trail is successful if the robot places an object into the target.

Door Opening The opening task (see Fig. 3c) requires the robot to approach a handle, grasp it, and then pull the door open. While conceptually simple, this task requires a great deal of precision, since the commanded actions must (almost) exactly match the direction of the door hinge. **Training:** We created a toy kitchen setting for this task, and collected 100 demonstrations of opening the door when starting from various initial positions. **Testing:** We evaluated on 20 real world test trials, where the robot’s initial position and the door’s initial state (e.g. position and hinge angle) were randomized. A trial was deemed successful if the door is opened fully.

Knob Turning In this final task (see Fig. 3d) the robot must approach a knob in the toy kitchen, and then turn it clock-wise. Similar to the opening task, knob turning requires very precise motor control, since the robot must grab the knob at exactly the right location and carefully rotate in order to turn. **Training:** We collect a dataset of 100 knob turning demonstrations. **Testing:** The method is evaluated on 20 trials, with randomized robot/knob initial positions. Success is determined by if the knob is turned far enough to cause an audible “click.”

4.2. Robot-Free Data Collection

A major strength of our setup is that training data (\mathcal{D}) can come from a different agent than the test time robot. We leverage this property to collect videos *directly* from humans operating a low-cost grabber stick (see Fig. 4, left), and use structure from motion [38] to recover actions. This allows us to collect demonstrations *without* an expensive tele-op setup [50, 23] or time-intensive kinesthetic demonstrations (i.e. manually moving the robot by hand). While based on prior work, our stick setup makes two important contributions: it is easier build and use than the highly instrumented setup in Song et. al. [42], and it provides higher




Figure 5: Trajectories executed on the robot using our learned distance function. For each task, we show the 1st person view (top) and 3rd person view images (bottom). We show the learned visual embedding can encode functional distances between states for challenging tasks, like pushing, pick and place, door opening, and knob turning.

quality images and action labels than the GoPro used in Young et. al. [48]. Please refer to Supplement A.2 for more information on our training and testing setups. We will release our dataset publicly.

5. Experiments

The following experiments seek to validate our distance learning framework on real control tasks (described above), and give some fundamental insights

as to why they work. First, our baseline study (see Sec. 5.1) compares our method against representative SOTA techniques in the field. In addition, in Sec. 5.2 we analyze a specific scenario with highly multi-modal actions, where our model especially shines. Finally, the ablation study (see Sec. 5.3) investigates which components of our system are actually needed for good performance.

| Task | Ours | BC [35, 48] | IBC [13] | IQL [20] |
|-----------------------|------------|-------------|----------|------------|
| <i>Pushing</i> | 85% | 70% | 70% | 65% |
| <i>Pick and Place</i> | 60% | 30% | 30% | 45% |
| <i>Door Opening</i> | 55% | 40% | 45% | 50% |
| <i>Knob Turning</i> | 40% | 20% | 20% | 35% |

Table 1: We compare success rates for our method versus the baselines on all four manipulation tasks. Our distance learning method outperforms a suite of representative robot learning baselines.

| # Demos | <i>Pushing</i> | | | <i>Pick and Place</i> | | |
|---------|----------------|-------------|-------------|-----------------------|------------|-----|
| | Ours | BC | IBC | Ours | BC | IBC |
| 50 | 70% | 50% | 40% | 0% | 0% | 0% |
| 100 | 85% | 70% | 70% | 10% | 10% | 0% |
| 200 | 90% | 80% | 100% | 20% | 20% | 10% |
| 400 | 100% | 100% | 100% | 60% | 30% | 30% |
| 600 | 100% | 100% | 100% | 70% | 50% | 40% |

Table 2: We compare our distance learning method versus the behavior cloning baselines (BC [35, 48] and IBC [13]) when trained on varying amounts of data. Distance learning is able to learn faster than the baselines and scale better with data.

5.1. Baseline Study

We first note that our paper uses only demonstration data and no online data. Therefore, we compare our method against three SOTA baselines – ranging from standard behavior cloning, to energy based modeling, and Offline Reinforcement Learning (RL) – that cover a wide range of prior LfD work. The baseline methods are described briefly below. To make the comparisons fair, we parameterize all neural networks with the same R3M representation backbone used by our method, and tune hyper-parameters for best possible performance. Please check Supplement B.2 for in depth details.

- **Behavior Cloning [35, 48] (BC):** BC learns a policy (via regression) that directly predicts actions from image observations: $\min_{\pi} \|\pi(I_t, I_g) - a_t\|_2$. This provides a strong comparison point for a whole class of LfD methods that focus on learning motor policies directly.
- **Implicit Behavior Cloning [13] (IBC):** IBC learns an energy based model that can predict actions during test time via optimization: $a_t = \operatorname{argmin}_a E(a, I_t, I_g)$. This method is conceptually very similar to behavior cloning, but has the potential to better handle multi-modal action distributions and discontinuous actions.
- **Implicit Q-Learning [20] (IQL):** IQL is an

offline-RL baseline that learns a Q function $Q(s, a) = Q((I_t, I_g), a_t)$, alongside a policy that maximizes it $\pi(I_t, I_g) = \operatorname{argmax}_a Q(s, a)$. Note that IQL’s training process require us to annotate our offline trajectories \mathcal{D} with a reward signal r_t for each time-step. While this can be done automatically in our setting (Supplement B.2), this is an additional assumption that could be very burdensome in the general case.

Our method is compared against the baselines on all four of our tasks (see Sec. 4.1). Results are reported in Table 1. Note how our distance learning approach achieves the highest success rates on all of four tasks. In addition, our method is much simpler than the strongest baseline (IQL), which requires an expert defined reward signal and uses a more complicated learning algorithm. The final test time performance is visualized in Fig. 5 and on our website: <https://agilabs.github.io/manipulate-by-seeing/>.

Data Scaling A common cited strength of BC (versus RL) is its ability to continuously improve with additional expert data. Can our method do the same? Table 2 compares our distance learning method versus the BC/IBC baselines (on pushing/pick-place tasks), when trained on various amounts of data. Note how our method both learns faster than BC/IBC, and continuously scales with added demonstrations. This suggests that distance learning is a powerful alternative to BC on robotic manipulation tasks, no matter how much training data is available.

5.2. Multi-Modality Experiment

A common challenge in policy learning is effectively handling multi-modal action predictions. We propose a simple test scenario that allows us to test our method in this setting without any other confounding variables. Specifically, we create a “obstacle pushing” task where the robot must push an object to the target without disturbing an obstacle placed in the middle (see Fig. 6). Observe that there are two equally valid ways that the robot could push the block; to the left or to the right. We collect a training dataset of 100 demos, equally split between going left and right. This creates action multi-modality where the agent must arbitrarily choose between two scenarios in order to successfully perform the task, but will fail if it simply averages the actions and pushes the block forward (into the obstacle). We expect that standard BC will fall into this failure mode, since it is trained using standard L2 action regression. In contrast, our method should still work since the low distance states will move the robot either to the left or right state, thus allowing it to make a discrete choice.




Figure 6: Our method can solve tasks with highly multi-modal action distributions that are difficult for the baselines. In this example, our distance learning controller successfully pushes the block around the obstacle, while Behavior Cloning learns to incorrectly push the block forward (i.e. predicts mean action).

To verify this hypothesis, we test our method, the BC baseline, and IQL (the strongest non-BC baseline) on 20 trials in this task setting. Success is judged similarly as to pushing (see Sec. 4.1), with the added condition that the robot should not push the obstacles into the target. Our method achieves a **95%** success rate, versus 0% for BC and 90% for IQL in this setting! As expected, our method is able to greatly improve upon BC in this setting, because it explicitly avoids action prediction. While IQL should not suffer as badly from multi-modality, thanks to its learned value function, our method is still able to outperform it despite being simpler algorithmically and trained without reward signal.

5.3. Ablation Study

Our final experiment removes components of our method in order to determine which parts were most crucial for control performance. The *representation ablation* replaces the R3M backbone in our neural network with a pre-trained ImageNet representation (using same ResNet-18 architecture for both). We expect the ImageNet weights to be weaker, since they were trained on a smaller dataset that did not include ego-centric videos (ImageNet vs Ego4D used by R3M). Additionally, the *dynamics ablation* removes the dynamics module during training, thus leaving the distance model/embeddings without any physical grounding. This is similar to the method from Hahn et. al. [15] (see

| | Ours | ImageNet Repr. | No Dynamics [15] |
|-----------------------|------------|----------------|------------------|
| <i>Pushing</i> | 90% | 40% | 60% |
| <i>Pick and Place</i> | 60% | 40% | 35% |
| <i>Door Opening</i> | 55% | 45% | 50% |
| <i>Knob Turning</i> | 40% | 15% | 40% |

Table 3: This table compares success rates for our method versus the ablations on the all four tasks. As you can see, removing the dynamics module from our method during training and training with a weaker representation both result in worse performance.

Sec. 2), with an added dynamics module. We expect this ablation to perform worse on manipulation tasks, because dynamics is more important in this setting v.s. visual navigation.

The two ablations are evaluated and compared against our method on all four test tasks (see Table 3). We find that our method strongly outperforms both of our baselines, which suggests that our intuition outline above is correct. In particular, we conclude that stronger representations can directly stronger control performance using our method. Additionally, this shows that the dynamics grounding is important to ensure that our learned distances transfer onto real robot hardware during test time, especially in manipulation tasks that involve object interactions/contact.

6. Discussion

This paper demonstrates that neural image representations can be more than just state representations: a simple metric defined within the embedding space can help infer robot actions. We leverage this insight to learn a distance function and dynamics function using minimal low-cost human data. These modules parameterize a robotic planner that is validated across 4 representative manipulation tasks. Despite its simplicity, our method is able to outperform standard imitation learning and offline-RL methods in the robot learning field. This is especially true in situations involving multi-modal action distributions, where our method entirely outclasses a standard BC baseline. Finally, the ablation study demonstrates that stronger representations directly result in stronger control performance, and that dynamics grounding is important for our scheme to work in practice.

Future Work We hope that this paper inspires future work in the space of representation learning and robotics. Follow up works should improve visual representations specifically for robotics, by more precisely capturing fine-grained interactions between gripper/-

hand and objects. This could improve performance on tasks like knob turning, where the pre-trained R3M encoder often struggled to detect small shifts of the gripper relative to the knob. We also hope that our contrastive learning method could be extended to learn entirely without action labels. This would allow us to train on large scale manipulation datasets (e.g. YouTube videos) and transfer the results directly to the control setting. Finally, it would be ideal if data collected using our low-cost stick (see Fig. 4) could be used with a stronger, more reliable (commercial) gripper, despite the domain gap. This would enable us to solve more tasks during test time.

References

- [1] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1, 2004.
- [2] OpenAI: Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, et al. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1):3–20, 2020.
- [3] Daniel Angelov, Yordan Hristov, Michael Burke, and Subramanian Ramamoorthy. Composing diverse policies for temporally extended tasks. *IEEE Robotics and Automation Letters*, 5(2):2658–2665, 2020.
- [4] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.
- [5] Yusuf Aytar, Tobias Pfaff, David Budden, Thomas Paine, Ziyu Wang, and Nando De Freitas. Playing hard exploration games by watching youtube. *Advances in neural information processing systems*, 31, 2018.
- [6] Shikhar Bahl, Abhinav Gupta, and Deepak Pathak. Human-to-robot imitation in the wild. *RSS*, 2022.
- [7] Aude Billard, Sylvain Calinon, Ruediger Dillmann, and Stefan Schaal. Survey: Robot programming by demonstration. *Handbook of robotics*, 59(BOOK_CHAP), 2008.
- [8] Annie S Chen, Suraj Nair, and Chelsea Finn. Learning generalizable robotic reward functions from” in-the-wild” human videos. *arXiv preprint arXiv:2103.16817*, 2021.
- [9] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021.
- [10] Sudeep Dasari and Abhinav Gupta. Transformers for one-shot imitation learning. In *CoRL 2020*, 2020.
- [11] Sudeep Dasari, Abhinav Gupta, and Vikash Kumar. Learning dexterous manipulation from exemplar object trajectories and pre-grasps. In *IEEE International Conference on Robotics and Automation 2023*, 2023.
- [12] Sudeep Dasari, Jianren Wang, Joyce Hong, Shikhar Bahl, Yixin Lin, Austin Wang, Abitha Thankaraj, Karanbir Chahal, Berk Calli, Saurabh Gupta, et al. Rb2: Robotic manipulation benchmarking with a twist. *arXiv preprint arXiv:2203.08098*, 2022.
- [13] Pete Florence, Corey Lynch, Andy Zeng, Oscar Ramirez, Ayzaan Wahid, Laura Downs, Adrian Wong, Johnny Lee, Igor Mordatch, and Jonathan Tompson. Implicit behavioral cloning. *Conference on Robot Learning (CoRL)*, 2021.
- [14] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of ego-centric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18995–19012, 2022.
- [15] Meera Hahn, Devendra Singh Chaplot, Shubham Tulsiani, Mustafa Mukadam, James M Rehg, and Abhinav Gupta. No rl, no simulation: Learning to navigate without navigating. *Advances in Neural Information Processing Systems*, 34:26661–26673, 2021.
- [16] Ankur Handa, Arthur Allshire, Viktor Makoviychuk, Aleksei Petrenko, Ritvik Singh, Jingzhou Liu, Denys Makoviichuk, Karl Van Wyk, Alexander Zhurkevich, Balakumar Sundaralingam, Yashraj Narang, Jean-Francois Lafleche, Dieter Fox, and Gavriel State. Dextreme: Transfer of agile in-hand manipulation from simulation to reality. *arXiv*, 2022.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [18] Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey Levine, and Chelsea Finn. Bc-z: Zero-shot task generalization with robotic imitation learning. In *Conference on Robot Learning*, pages 991–1002. PMLR, 2022.
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [20] Ilya Kostrikov, Ashvin Nair, and Sergey Levine. Offline reinforcement learning with implicit q-learning. *arXiv preprint arXiv:2110.06169*, 2021.
- [21] Aviral Kumar, Rishabh Agarwal, Dibya Ghosh, and Sergey Levine. Implicit under-parameterization inhibits data-efficient deep reinforcement learning. *arXiv preprint arXiv:2010.14498*, 2020.
- [22] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. Rma: Rapid motor adaptation for legged robots. *arXiv preprint arXiv:2107.04034*, 2021.
- [23] Vikash Kumar and Emanuel Todorov. Mujoco haptix: A virtual reality system for hand manipulation. In

- Humanoid Robots (Humanoids)*, 2015 IEEE-RAS 15th International Conference on, pages 657–663. IEEE, 2015.
- [24] Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning quadrupedal locomotion over challenging terrain. *Science robotics*, 5(47):eabc5986, 2020.
- [25] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- [26] Yixin Lin and Austin Wang. Polymetis: a real-time pytorch controller manager. In <https://github.com/facebookresearch/polymetis>, 2021.
- [27] Corey Lynch, Mohi Khansari, Ted Xiao, Vikash Kumar, Jonathan Tompson, Sergey Levine, and Pierre Sermanet. Learning latent plans from play. In *Conference on robot learning*, pages 1113–1132. PMLR, 2020.
- [28] Andriy Mnih and Koray Kavukcuoglu. Learning word embeddings efficiently with noise-contrastive estimation. *Advances in neural information processing systems*, 26, 2013.
- [29] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3m: A universal visual representation for robot manipulation. *CoRL*, 2022.
- [30] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [31] Simone Parisi, Aravind Rajeswaran, Senthil Purushwalkam, and Abhinav Gupta. The unsurprising effectiveness of pre-trained vision models for control. *arXiv preprint arXiv:2203.03580*, 2022.
- [32] Ilija Radosavovic, Tete Xiao, Stephen James, Pieter Abbeel, Jitendra Malik, and Trevor Darrell. Real-world robot learning with masked visual pre-training. *CoRL*, 2022.
- [33] Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel Todorov, and Sergey Levine. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. *arXiv preprint arXiv:1709.10087*, 2017.
- [34] Frank Rosenblatt. Principles of neurodynamics. perceptrons and the theory of brain mechanisms. Technical report, Cornell Aeronautical Lab Inc Buffalo NY, 1961.
- [35] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635, 2011.
- [36] Stefan Schaal. Is imitation learning the route to humanoid robots? *Trends in cognitive sciences*, 3(6):233–242, 1999.
- [37] Karl Schmeckpeper, Oleh Rybkin, Kostas Daniilidis, Sergey Levine, and Chelsea Finn. Reinforcement learning with videos: Combining offline observations with interaction. In *Conference on Robot Learning*, pages 339–354. PMLR, 2021.
- [38] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [39] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, Sergey Levine, and Google Brain. Time-contrastive networks: Self-supervised learning from video. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 1134–1141. IEEE, 2018.
- [40] Pierre Sermanet, Kelvin Xu, and Sergey Levine. Un-supervised perceptual rewards for imitation learning. *arXiv preprint arXiv:1612.06699*, 2016.
- [41] Nur Muhammad Mahi Shafiullah, Zichen Jeff Cui, Ariuntuya Altanzaya, and Lerrel Pinto. Behavior transformers: Cloning k modes with one stone. In *Thirty-Sixth Conference on Neural Information Processing Systems*, 2022.
- [42] Shuran Song, Andy Zeng, Johnny Lee, and Thomas Funkhouser. Grasping in the wild: Learning 6dof closed-loop grasping from low-cost demonstrations. *Robotics and Automation Letters*, 2020.
- [43] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [44] Michita Imai Takuma Seno. d3rlpy: An offline deep reinforcement learning library. In *NeurIPS 2021 Offline Reinforcement Learning Workshop*, December 2021.
- [45] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pages 5026–5033. IEEE, 2012.
- [46] Faraz Torabi, Garrett Warnell, and Peter Stone. Generative adversarial imitation from observation. *arXiv preprint arXiv:1807.06158*, 2018.
- [47] Annie Xie, Avi Singh, Sergey Levine, and Chelsea Finn. Few-shot goal inference for visuomotor learning and planning. In *Conference on Robot Learning*, pages 40–52. PMLR, 2018.
- [48] Sarah Young, Dhiraj Gandhi, Shubham Tulsiani, Abhinav Gupta, Pieter Abbeel, and Lerrel Pinto. Visual imitation made easy. In *Conference on Robot Learning*, pages 1992–2005. PMLR, 2021.
- [49] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning (CoRL)*, 2019.
- [50] Tianhao Zhang, Zoe McCarthy, Owen Jow, Dennis Lee, Xi Chen, Ken Goldberg, and Pieter Abbeel. Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5628–5635. IEEE, 2018.