

Symmetric Models for Visual Force Policy Learning

Colin Kohler¹, Anuj Shrivatsav Srikanth¹, Eshan Arora¹, Robert Platt¹

Abstract—While it is generally acknowledged that force feedback is beneficial to robotic control, applications of policy learning to robotic manipulation typically only leverage visual feedback. Recently, symmetric neural models have been used to significantly improve the sample efficiency and performance of policy learning across a variety of robotic manipulation domains. This paper explores an application of symmetric policy learning to visual-force problems. We present Symmetric Visual Force Learning (SVFL), a novel method for robotic control which leverages visual and force feedback. We demonstrate that SVFL can significantly outperform state of the art baselines for visual force learning and report several interesting empirical findings related to the utility of learning force feedback control policies in both general manipulation tasks and scenarios with low visual acuity.

I. INTRODUCTION

There are a variety of manipulation tasks where it is essential to use both vision and force feedback as part of the control policy. Peg insertion with tight tolerances, for example, is a task that is nearly impossible to solve without leveraging force feedback in some form. The classical approach is to use an admittance controller with a remote center of compliance to help the peg slide into the hole [1]. However, this is a very limited use of force feedback and it seems like it should be possible to use force information in a more comprehensive way. One of the core challenges is the difficulty in simulating the complex force interactions that happen at the robot end effector which primarily depend upon the contact-modelling utilized by the physics engine. While there have been major efforts to improve contact simulations by refining the contact geometry, friction model, and contact constraints [2], [3], [4], state-of-the-art physics engines often violate real-world physical constraints limiting the applicability of simulation-based models to the real world.

An obvious alternative approach is to leverage machine learning, i.e. model free reinforcement learning (RL), to obtain force feedback assisted policies. This is in contrast to vision-only RL where the policy only takes visual feedback [5], [6], [7], [8]. In *visual force* RL, there is the possibility to adapt control policies directly to the mechanical characteristics of the system as they exist in the physical world, without the need to model those dynamics first. However, this assumes that we can run RL online directly in the physical world, something that is hard to do due to the poor sample efficiency of RL. RL is well known to require an enormous amount of data in order to learn even simple policies effectively. While visual force RL might, in

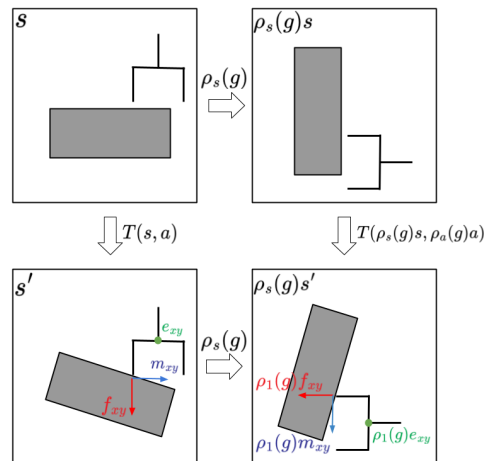


Fig. 1: **Symmetries in Visual-Force Domains.** An example of the $O(2)$ symmetries in the visual-force policy learning problem. Note that the same rotational symmetries which apply to the image s also apply to the xy components of the force (f_{xy} , m_{xy}) and end effector state (e_{xy}).

principle, be able to learn effective policies, this sample inefficiency prevents us from learning policies directly on physical equipment. In order to improve the sample efficiency of RL in visual force problems, one common approach is to learn a helpful latent representation during a pretraining phase [9], [10], [11], [12]. This generally takes the form of self-supervised robot “play” in the domain of interest that must precede actual policy learning. Unfortunately, this is both cumbersome and brittle as the latent representation does not generalize well outside the situations experienced during the play phase. This is especially prevalent in the visual force domain as the noisy nature of force sensors means there will be many force observations not experienced during pretraining leading to poor latent predictions during policy learning.

This paper develops an alternative approach to the problem of visual force learning based on exploiting domain symmetries using equivariant learning [13]. Recently, symmetric neural networks have been shown to dramatically improve the sample efficiency of RL in robotic manipulation domains [14], [15]. However, this work has focused exclusively on visual feedback and has not yet been applied to visual force learning. This is of particular note as the domain symmetries used by these equivariant models are also present in both force and proprioceptive data. This can be seen in Fig. 1 where the image, force, and proprioceptive observations are all invariant under rotations in the xy plane. This paper makes three main contributions. First, we propose a novel method for visual force policy learning called Symmetric

¹Northeastern University, Boston, MA 02115, USA.

Visual Force Learning (SVFL) which exploits the underlying symmetry of manipulation tasks to improve sample efficiency and performance. Second, we empirically evaluate the importance of force feedback assisted control across a variety of manipulation domains and find that force feedback is helpful for nearly all tasks, not just contact-rich domains like peg insertion where we would expect it to be important. Finally, we explore the role of force-assisted policies in domains with low visual acuity and characterize the degree to which force models can compensate for poor visual information.

II. RELATED WORK

Contact-Rich Manipulation. Contact-rich manipulation tasks, i.e. peg insertion, screw fastening, edge following, etc., are well-studied areas in robotic manipulation due to their prevalence in manufacturing domains. These tasks often are solved by hand-engineered policies which utilize force feedback and very accurate state estimation [1], resulting in policies that perform well in structured environments but do not generalize to the large range of real-world variability. More recent work has proposed the use of reinforcement learning to address this by training neural network policies which combine vision and proprioception [6], [16], [17]. However, while these methods have been shown to perform well across a variety of domains and task variations, they require a high level of visual acuity, such that the task is solvable solely using image observations. In practice, this means these methods are unsuitable for a large portion of contact-rich manipulation tasks which require a high degree of precision and often include visual obstructions.

Multimodal Learning. A common approach to multimodal learning is to first learn a latent dynamics model which compactly represents the high-dimensional observations and then use this model for policy learning. This technique has recently been adapted for use in various robotics domains to combine various types of heterogeneous data. [18] combine vision and haptic information using a GAN but do not utilize their latent representation for manipulation policies. [12] and [11] learn physics models from cross-modal visual-tactile data for a series of tasks but they do not use this learned representation for either a hand-crafted policy or policy learning. Our work is most closely related to [10], [9] which we use as baselines in this work. [10] combine vision, force, and proprioceptive data using a variational latent model learned from self-supervision and use this model to learn a policy for peg insertion. [9] learn a multimodal latent heatmap using a cross-modal visuo-tactile transformer (VTT) which distributes attention spatially. They show that by combining VTT with stochastic latent actor critic (SLAC), they can learn policies that can solve a number of manipulation tasks. In comparison to these works, we propose a sample-efficient deterministic multimodal representation that is learned end-to-end without the need for pretraining.

Equivariant Neural Networks. Equivariant networks were first introduced as G-Convolutions [19] and Steerable CNNs [13], [20], [21]. Since their inception they have been applied across varied datatypes including images [20],

spherical data [22], [23], and point clouds [24]. More recent work has expanded the use of equivariant networks to reinforcement learning [15], [7], [25] and robotics [26], [27], [28], [29]. Compared to these prior works which focus on a single data modality, this work studies the effectiveness of combining various heterogeneous datatypes while preserving the symmetry inherent in each of these data modalities.

III. BACKGROUND

Equivariant Neural Networks. A function is equivariant if it respects the symmetries of its input and output spaces. Specifically, a function $f : X \rightarrow Y$ is *equivariant* with respect to a symmetry group G if it commutes with all transformations $g \in G$, $f(\rho_x(g)x) = \rho_y(g)f(x)$, where ρ_x and ρ_y are the *representations* of the group G that define how the group element $g \in G$ acts on $x \in X$ and $y \in Y$, respectively. An equivariant function is a mathematical way of expressing that f is symmetric with respect to G : if we evaluate f for various transformed versions of the same input, we should obtain transformed versions of the same output. Although this symmetry can be learned [30], in this work we require the symmetry group G and representation ρ_x to be known at design time. For example, in a convolutional model, this can be accomplished by tying the kernel weights together to satisfy $K(gy) = \rho_{out}(g)K(y)\rho_{in}(g)^{-1}$, where ρ_{in} and ρ_{out} denote the representation of the group operator at the input and output of the layer [31]. End-to-end equivariant models can be constructed by combining equivariant convolutional layers and equivariant activation functions. In order to leverage symmetry in this way, it is common to transform the input so that standard group representations work correctly, e.g., to transform an image to a top-down view so that image rotations correspond to object rotations.

Extrinsic Equivariance. Often real-world problems contain symmetry corruptions such as oblique viewing angles and occlusions. This is particularly prevalent in robotics domains where the state of the world is rarely fully observable. In these domains we consider the symmetry to be *latent* where we know that some symmetry is present in the problem but cannot easily express how that symmetry acts in the input space. We refer to this relationship as *extrinsic equivariance* [25], where the equivariant constraint in the equivariant network enforces equivariance to out-of-distribution data. While extrinsic equivariance is not ideal, it does not necessarily increase error and has been shown to provide significant performance improvements in reinforcement learning [25].

IV. APPROACH

A. Problem Statement

We model the visual force control problem as a discrete time finite horizon Markov decision process (MDP), $\mathcal{M} = (S, A, \mathcal{T}, R, \gamma)$, where states $s \in S$ encode visual, force, and proprioceptive data and actions $a \in A$ command small end effector displacements. This MDP transitions at a frequency of 20 Hz and the commanded hand displacements are provided as positional inputs to a lower level Cartesian

space admittance controller that runs at 500Hz with a fixed stiffness. The hand is constrained to point straight down at the table (along the $-z$ direction).

State is a tuple $s = (I, f, e) \in S$. $I \in \mathbb{R}^{4 \times h \times w}$ is a 4-channel RGB-D image captured from a fixed camera pointed at the workspace. $f = (f_{xy}, f_z, m_{xy}, m_z) \in \mathbb{R}^{T \times 6}$ is a $T \times 6$ time series of the last T measurements from a six-axis wrist force-torque sensor transformed into the robot base frame. $e = (e_\lambda, e_{xy}, e_z, e_\theta) \in \mathbb{R}^5$ is the configuration of the end effector where $e_\lambda \in E_\lambda$ is the hand open width, (e_{xy}, e_z) are the Cartesian coordinates of the hand, and e_θ is the orientation of the hand about the $-z$ axis. Actions are represented by $a = (\lambda, \Delta p) \in A \subseteq \mathbb{R}^5$ where $\lambda \in \mathbb{R}$ is the desired gripper open width and $\Delta p = (\Delta p_{xy}, \Delta p_z, \Delta p_\theta) \in \mathbb{R}^4$ is the desired delta pose of the gripper with respect to the current pose p .

B. $O(2)$ Symmetries in Visual Force Domains

In order to leverage symmetric models for visual force policy learning, we utilize the group invariant MDP framework. A group invariant MDP is an MDP with reward and transition functions that are invariant under the group action, $R(s, a) = R(\rho_s(g)s, \rho_a(g)a)$ and $\mathcal{T}(s, a, s') = \mathcal{T}(\rho_s(g)s, \rho_a(g)a, \rho_s(g)s')$, for elements of an appropriate symmetry group $g \in G$ [14], [32]. ρ_s and ρ_a are representations of the group G that define how group elements act on state and action. This paper focuses on discrete subgroups of $O(2)$ such as the dihedral groups D_4 or D_8 that represent rotations and reflections in the xy plane, i.e. the plane of the table. We utilize the D_8 group in our experiments.

In order to express visual force manipulation as a group invariant MDP, we must define how the group operates on state and action such that the transition and reward invariance equalities described above are approximately satisfied. State is $s = (I, f, e) = (I, f_{xy}, f_z, m_{xy}, m_z, e_{xy}, e_z, e_\lambda)$. Since we are focused on rotations and reflections in the plane about the z axis, only the xy variables are affected. Therefore, the group $g \in SO(2)$ acts on s via

$$\begin{aligned} \rho_s(g)I &= \rho_0(g)I \\ \rho_s(g)f &= (\rho_1(g)f_{xy}, f_z, \rho_1(g)m_{xy}, m_z) \\ \rho_s(g)e &= (\rho_1(g)e_{xy}, e_z, e_\lambda) \end{aligned}$$

where $\rho_0(g)$ is a linear operator that rotates/reflects the pixels in an image by g and $\rho_1(g)$ is the standard representation of rotation/reflection in the form of a 2×2 orthonormal matrix. Turning to action, $a = (\lambda, \Delta p_{xy}, \Delta p_z, \Delta p_\theta)$, we define $\rho_a(g)a = (\lambda, \rho_1(g)\Delta p_{xy}, \Delta p_z, \Delta p_\theta)$. Given these definitions, visual force manipulation satisfies the transition and reward invariance constraints, $R(s, a) = R(\rho_s(g)s, \rho_a(g)a)$ and $\mathcal{T}(s, a, s') = \mathcal{T}(\rho_s(g)s, \rho_a(g)a, \rho_s(g)s')$. This is illustrated for transition invariance in Fig. 1.

C. Model Architecture

As we discuss in the next section, we do policy learning using SAC which requires a critic (a Q -function) and an actor. In our method, both actor and critic employ the

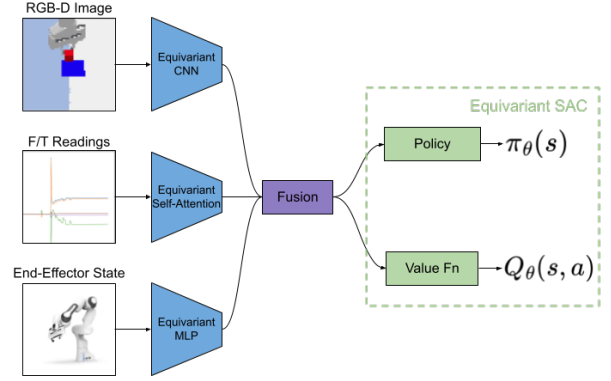


Fig. 2: **Symmetric Visual Force Learning.** SVFL combines visual, force, and proprioceptive feedback into a single joint-representation while respecting the symmetry inherent to each data modality.

same encoder architecture which encodes state into a latent representation. Since our state $s = (I, f, e) \in S$ is multimodal (i.e. vision, force, and proprioception) our backbone is actually three encoders, the output of which is concatenated (Fig. 2). The image encoder (top left in Fig. 2) is a series of seven equivariant convolutional layers. The force encoder (middle left) is a single equivariant self-attention layer. The proprioceptive encoder (bottom left) is a four-layer equivariant MLP. In each of these encoders, the model respects the equivariance and invariance of each data modality corresponding to the relationships described in Section IV-B. These equivariant networks are implemented using the escnn [33], [34] library, where all the hidden layers are defined using regular representations.

The force encoder is of particular note due to its use of single-headed self-attention. The input is a set of T tokens, $f \in \mathbb{R}^{T \times 6}$, that encode the most recent T measurements from the force-torque sensor. In order to make this model equivariant, we convert each of the key, query, and value networks to become equivariant models. For the standard implementation of self-attention, $\text{Attn} = \text{softmax}(fW^Q(fW^K)^T)fW^V$, the resulting group self attention operation is equivariant [35]:

$$\begin{aligned} \text{Attn}(X_f\Gamma) &= \text{softmax}(X_f\Gamma W^Q(X_f\Gamma W^K)^T)X_f\Gamma W^V \\ &= \text{softmax}(X_f W^Q\Gamma(X_f W^K\Gamma)^T)X_f W^V\Gamma \\ &= \text{softmax}(X_f W^Q\Gamma\Gamma^T(X_f W^K)^T)X_f W^V\Gamma \\ &= \text{softmax}(X_f W^Q(X_f W^K)^T)X_f W^V\Gamma \\ &= \text{Attn}(X_f)\Gamma, \end{aligned}$$

where, for simplicity of this analysis, we define Γ to be the linear representation of the action of a group element $g \in G$ and $X_f \in \mathbb{R}^{T \times 6 \times |G|}$.¹

D. Equivariant SAC

For policy learning, we use Soft Actor Critic (SAC) [36] combined with the model architecture described above. This can be viewed as a variation of Equivariant SAC [14] that is adapted to visual force control problems. The policy is

¹Although we omit the positional encoding here, this does not affect the result [35].

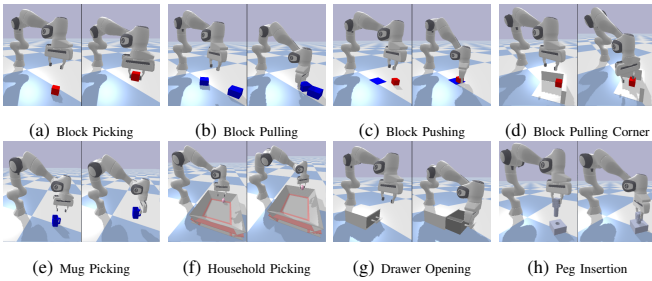


Fig. 3: **Tasks.** The manipulation domains from the BulletArm benchmark [38] implemented in PyBullet [39]. (Left) Initial state. (Right) Goal state.

TABLE I: Number of trainable parameters in the policy learning tasks. Due to being latent representation learning methods, VTT and PoE utilize shared encoders between the actor and the critic so that the number of parameters is smaller than SVFL and CNN. Additionally, we utilize a smaller model for PoE as increasing the size of the PoE model has been shown to worsen performance [9].

Network	SVFL	CNN	VTT	POE
# of Parameters	2.4E6	2.5E6	1.19E6	2.9E5

a network $\pi : S \rightarrow A \times A_\sigma$, where A_σ is the space of action standard deviations. We define the group action on the action space of the policy network $\bar{a} \in A \times A_\sigma$ as: $\rho_{\bar{a}}(g)\bar{a} = (\rho_a(g)a, a_\sigma)$, where $a_\sigma \in A_\sigma$ and $g \in G$. The actor network π is defined as a mapping $s \mapsto \bar{a}$ that satisfies the following equivariance constraint: $\pi(\rho_s(g)s) = \rho_a(g)(\pi(s))$. The critic is a Double Q-network [37]: $q : S \times A \rightarrow \mathcal{R}$ that satisfies an invariant constraint: $q(\rho_s(g)s, \rho_a(g)a) = q(s, a)$. Both the actor and critic are implemented using the escnn library. For the critic, the output is a trivial representation. For the actor, the output is a mixed representation containing one standard representation for the (x, y) actions, one signed representation for the θ action, and seven trivial representations for the (λ, z) actions alongside the action standard deviations.

V. EXPERIMENTS

We performed a series of experiments both in simulation and on physical hardware to validate our approach, Symmetric Visual Force Learning (SVFL). First, we benchmark SVFL’s performance in simulation against several alternative approaches. Second, we perform ablations that measure the contributions of different input modalities for different tasks under both ideal and degraded visual observations. Finally, we validate the approach on physical hardware.

A. Simulated Experiments

Tasks. We evaluate SVFL across nine manipulation tasks from the BulletArm benchmark [38] which uses the PyBullet [39] simulator: Block Picking, Block Pushing, Block Pulling, Block Corner Pulling, Mug Picking, Household Picking, Peg Insertion, Drawer Opening, and Drawer Closing (Fig. 3). The workspace’s size is $0.4m \times 0.4m \times 0.26m$. The minimum z height is slightly beneath the table allowing the arm to come in contact with the table. The pixel size of the visual observation is $4 \times 76 \times 76$ and is cropped to $4 \times 64 \times 64$ during training and testing. The force data consists of the most recent 64 readings from the F/T sensor. The maximum

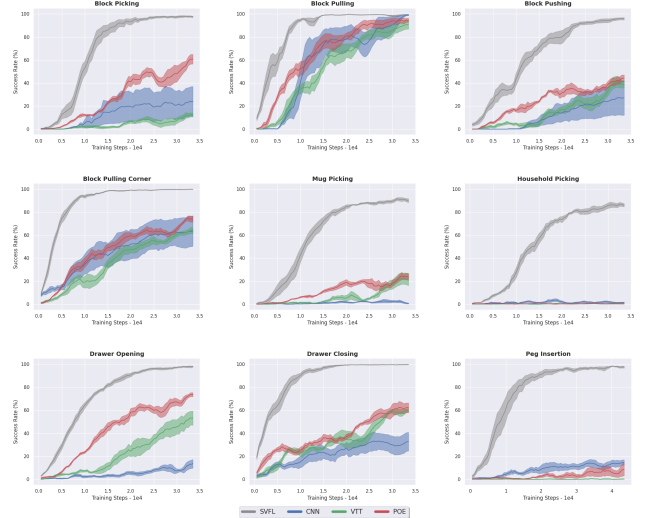


Fig. 4: **Baseline Comparison.** Comparison of SVFL (gray) with baselines. Greedy evaluation policy is shown in terms of success rate. In all of our experiments, results are averaged over 5 random seeds and the evaluation is performed every 500 training steps. Shading denotes standard error.

movement allowed for any action is limited to $\Delta x, \Delta y, \Delta z \in [-2.5cm, 2.5cm]$, $\Delta \theta \in [-\frac{\pi}{16}, \frac{\pi}{16}]$, $\lambda \in [e_{min}, e_{max}]$ where e_{min} and e_{max} are the joint limits of the gripper. For all tasks, a sparse reward function is used where a reward of +1 is given at task completion and 0 otherwise.

Baselines. We benchmark our method against two prominent alternative methods for visual force (or visual tactile) learning that have been proposed recently: Visuo-Tactile Transformers (VTT) [9] and Product of Experts (PoE) [10]. We also compare against a non-symmetric version of our model that is the same in every way except that it does not use equivariant layers (CNN). Both PoE and VTT are latent representation methods which rely on a self-supervised pretraining phase to build a compact latent representation of the underlying states providing increased sample efficiency. Due to this pretraining, these methods represent attractive options for on-robot policy learning. In both baselines we used the encoder architectures proposed in [9] which were shown to outperform those in [10]. PoE encodes the different input modalities independently using separate encoders and combines them using product of experts [10]. VTT combines modalities by using cross-modal attention on force to build latent representations that focus attention on important task features in the visual state space [9]. For further details on these baselines, see [9], [10]. The latent encoders are pretrained for 10^4 steps on expert data to predict the reconstruction of the state, contact and alignment embeddings, and the reward. All methods use Prioritized Experience Relay (PER) [40] pre-loaded with 50 episodes of the expert data. We augment the transitions by randomly cropping the visual observations and applying random rotations to the full observation. Parameter counts for all method can be found in Table I.

Results. We compared our method (SVFL) against the two baselines (POE and VTT) and the non-symmetric model (CNN) on the nine domains described above. Results are

shown in Fig. 4. All results are averaged over five runs starting from independent random seeds. When compared to the baselines, SVFL has significantly higher success rates and sample efficiency in all cases. We note the lower performance of VTT and PoE compared to [9] (namely in the Peg Insertion task) which we attribute to our use of sparse-reward functions which has been shown to lead to poor latent representations when using SLAC [41].

B. Sensor Modality Ablation

Although it is intuitive that force data should help learn better policies on manipulation tasks, especially on contact rich tasks like peg insertion, it is important to validate this assumption and to measure the benefits that can be gained by using both vision and force feedback rather than vision alone. Recall that our state representation can be factored into three modalities, $s = (I, f, e)$, where I is an image (vision), f is force, and e is the configuration of the robot hand (proprioception). Here, we compare the performance of SVFL with all three modalities against a vision-only model, a vision/force model, and a vision/proprioception model on the same tasks as in Section V-A. Results are shown in Fig. 5. The results indicate that the inclusion of each additional sensor modality improves sample efficiency and performance for policy learning with all three sensor modalities performing best in most cases. However, notice that the degree to which force (and proprioceptive) data helps depends upon the task. For example, the addition of force feedback drastically improves performance in Peg Insertion but has almost no effect in Block Pulling. There are, however, many tasks between these extremes. In Drawer Opening and Block Picking the force-aware policy converges to a slightly higher success rate than the non-force assisted policies. The fact that force feedback is usually helpful, even in tasks where one might not expect it, is interesting. This suggests that there is real value in incorporating force feedback into a robotic learning pipeline, even when there is a non-trivial cost to doing so.

C. Role of Force Feedback When Visual Acuity is Degraded

We also perform experiments in the context of degraded visual acuity to determine what happens if the visual input to our model is scaled down significantly. Specifically, we evaluate the model on RGB-D images rescaled (bilinear interpolation) to four different sizes: 64×64 , 32×32 , 16×16 , and 8×8 . Aside from the rescaling, all other aspects of the model match the SVFL method detailed in the previous section. This experiment gives an indication of how force data can compensate for low resolution cameras, cloudy environments, or smudged camera lenses. Fig. 6 shows performance at convergence at the four different levels of visual resolution. We note several interesting observations. First, the importance of visual acuity is dependant on the task, e.g. high visual acuity is very important for Block Picking but not very important for Block Pulling. Second, force information generally tends to help the most in low visual acuity scenarios. Finally, while force data generally

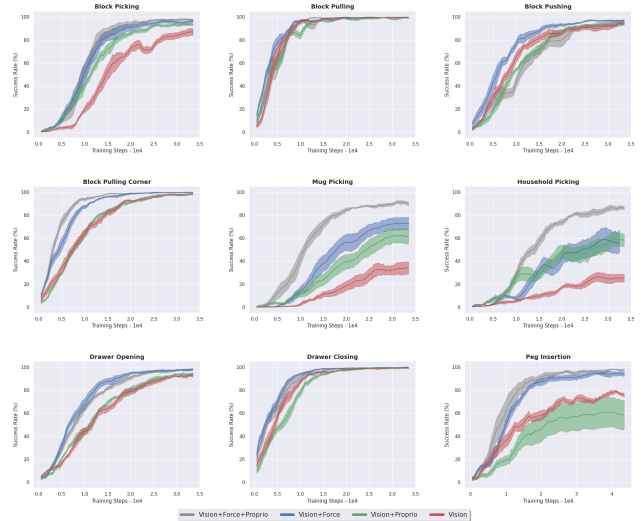


Fig. 5: **Sensor Modality Ablation.** Comparison of the full SVFL model (gray) to subsets of the data modalities.

improves performance, it cannot compensate for the loss of information in extreme visual degradation in tasks which require high visual acuity.

D. Real-World On-Robot Policy Learning

We repeat the simulated Block Picking policy learning experiment from Section V-B in the real world to evaluate our methods performance in the real-world. Fig. 7 shows the experimental setup which includes a UR5e robotic arm, a Robotiq Gripper, a wrist-mounted force-torque sensor, and a Intel RealSense camera. The block is a $5cm$ wooden cube that is randomly posed in the workspace. We utilize AprilTags to track the block for use in reward/termination checking and to automatically reset the workspace by moving the block to a new pose at the start of each episode. These tags are not utilized during policy learning. In order to facilitate faster learning, we modify a number of environmental parameters in our real-world setup. We use a workspace size as of $0.3m \times 0.3m \times 0.26m$ and a sparse reward function. We increase the number of expert demonstrations to 100 (from 50) and reduce the maximum number of steps per episode to 25 (from 50). Additionally, we reduce the action space by removing control of the gripper orientation and increase the maximum amount of movement the policy can take in one step to $5cm$ (from $2.5cm$). We utilize the same model architecture as in Section V-A.

Fig. 7 shows the learning curve of the full SVFL model alongside the various subsets of data modalities available to our method. We train all models for 3000 steps taking around 4 hours. As in the simulation results, the full SVFL model is both more sample efficient and outperforms SVFL modality subsets. Additionally, we see that force sensing is a vital component in this setting with the force-aware models achieving a 90% success rate compared to the 60% success rate of the non-force aware models (at 3000 training steps).

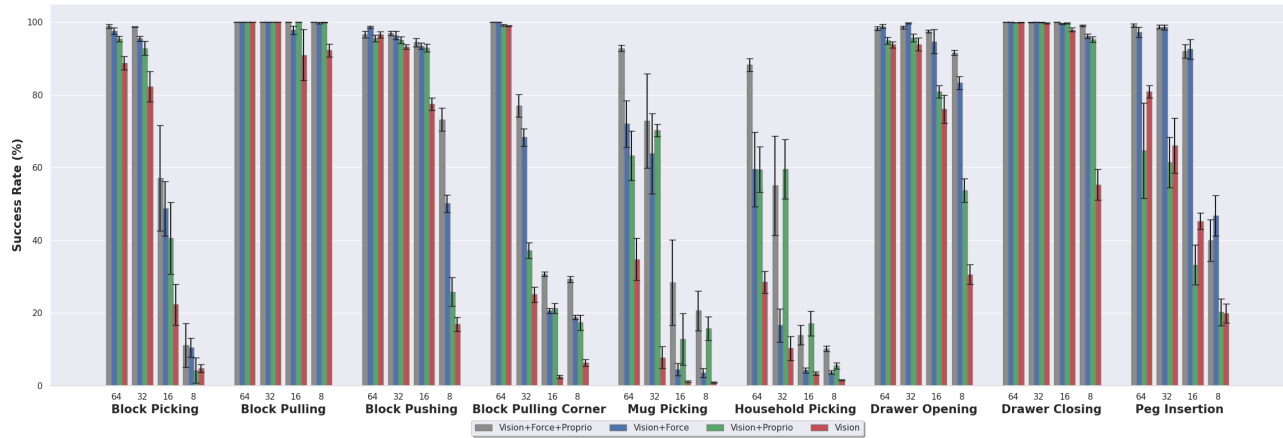


Fig. 6: **Performance Under Degraded Visual Acuity.** Comparison of the full SVFL model (gray) versus SVFL with subsets of the data modalities under visual acuity degradation. Performance is given after all models are trained to convergence.

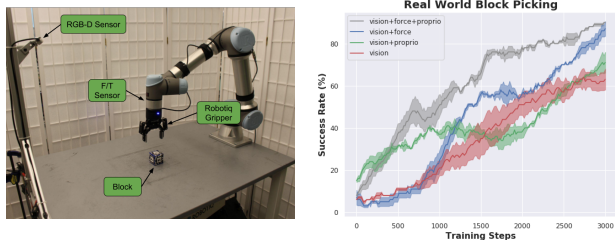


Fig. 7: **On-Robot Policy Learning.** (Left) Robotic setup. (Right) Comparison of the full SVFL model (gray) versus SVFL with subsets of the data modalities in the real-world on the Block Picking task. Results are averaged over 3 runs.

E. Real-World Block Centering

To further examine the advantages of SVFL, we conduct a supervised learning experiment using the same experimental setup in Fig. 7. Here we learn a function, $h : (I, f) \mapsto \{0, 1\}^4$, that maps visual-force observations to a four-way classification denoting the direction in which the gripper would need to move in order to grasp the block after a finger collides with the block. The idea was to mimic the most common failure case we see during policy learning in block picking where the grasp was slightly offset from the block. The dataset is generated by a human teleoperator where each sample is the most recent sensor observations immediately following the collision. The goal of the teleoperator was to mimic a failed grasp where one finger came into contact with the block. We generate 200 data samples and split the dataset into 100 training samples and 100 testing samples. We generated a diverse set of interactions by varying the position of the gripper in relation to the block, the amount of force (by varying the amount of movement when coming into contact with the block), and the pose of the block.

We compare the classification accuracy of the baseline SVFL model against the non-symmetric version of the model. We examine the effect of three different types of input: Vision Only (V), Force Only (F), and Vision & Force (V+F). In each case, in order to measure the models’ ability to generalize, we evaluated the performance on training sets of differing sizes including 10, 25, 50, and 100 samples. Table II shows the accuracy of the models on the held-out

TABLE II: **Experiment on Robotic Hardware.** Prediction accuracy (%) on the test set for models trained with different amounts of training data. We compare the performance of equivariant and non-symmetric versions of the vision encoder (V), the force encoder (F), and the fusion of these two encoders (V+F). Mean and standard error is given over three runs.

		# of Training Samples			
Encoder		10	25	50	100
V	CNN	38.7 ± 1	41.4 ± 3	35.1 ± 7	59.5 ± 6
	SVFL	51.3 ± 3	59.5 ± 5	67.7 ± 4	94.6 ± 4
F	CNN	21.6 ± 1	36.0 ± 3	42.3 ± 1	45.9 ± 7
	SVFL	30.6 ± 3	54.5 ± 2	81.1 ± 4	92.8 ± 1
V+F	CNN	38.7 ± 3	38.7 ± 6	45.0 ± 3	73.9 ± 3
	SVFL	57.8 ± 4	63.2 ± 3	87.6 ± 9	98.4 ± 1

test dataset. Notice that in all cases, the symmetric model does much better than its non-symmetric counterpart, both for differently sized training sets as well as input types.

VI. DISCUSSION & LIMITATIONS

This paper proposes Symmetric Visual Force Learning (SVFL), an approach to policy learning with visual force feedback that incorporates $SE(2)$ symmetries into the model. Our experiments demonstrate that SVFL outperforms two recent high profile benchmarks, PoE [10] and VTT [9], by a significant margin both in terms of learning speed and final performance. We also report a couple of interesting empirical findings. First, we find that force feedback is helpful across a wide variety of policy learning scenarios, not just those where one would expect force feedback to help, i.e. Peg Insertion. Second, we find that the positive effect of incorporating force feedback increases as visual acuity decreases. A limitation of this work is that although we expect that our framework is extensible to haptic feedback, this paper focuses on force feedback only. Additionally, we constrain our problem to top-down manipulation and planar symmetries in $SE(2)$ and therefore there is significant scope to extend this to $SE(3)$ symmetries. Finally, this paper focuses primarily on RL but the encoder architectures should be widely applicable to other learning techniques such as imitation learning.

REFERENCES

- [1] D. E. Whitney, "Historical perspective and state of the art in robot force control," *The International Journal of Robotics Research*, vol. 6, no. 1, pp. 3–14, 1987.
- [2] R. Tedrake and the Drake Development Team, "Drake: Model-based design and verification for robotics," 2019. [Online]. Available: <https://drake.mit.edu>
- [3] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 5026–5033.
- [4] J. Liang, V. Makovychuk, A. Handa, N. Chentanez, M. Macklin, and D. Fox, "Gpu-accelerated robotic simulation for distributed reinforcement learning," in *Conference on Robot Learning*. PMLR, 2018, pp. 270–282.
- [5] R. Platt, C. Kohler, and M. Gualtieri, "Deictic image mapping: An abstraction for learning pose invariant manipulation policies," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 8042–8049.
- [6] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016.
- [7] D. Wang, C. Kohler, and R. Platt, "Policy learning in se (3) action spaces," in *Conference on Robot Learning*. PMLR, 2021, pp. 1481–1497.
- [8] C. Kohler and R. Platt, "Visual foresight with a local dynamics model," in *The International Symposium of Robotics Research*. Springer, 2022, pp. 67–82.
- [9] Y. Chen, M. V. der Merwe, A. Sipos, and N. Fazeli, "Visuo-tactile transformers for manipulation," in *6th Annual Conference on Robot Learning*, 2022. [Online]. Available: <https://openreview.net/forum?id=JqqSTgdQ85F>
- [10] M. A. Lee, Y. Zhu, P. Zachares, M. Tan, K. Srinivasan, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg, "Making sense of vision and touch: Learning multimodal representations for contact-rich tasks," *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 582–596, 2020.
- [11] W. Zheng, H. Liu, and F. Sun, "Lifelong visual-tactile cross-modal learning for robotic material perception," *IEEE transactions on neural networks and learning systems*, vol. 32, no. 3, pp. 1192–1203, 2020.
- [12] N. Fazeli, M. Oller, J. Wu, Z. Wu, J. B. Tenenbaum, and A. Rodriguez, "See, feel, act: Hierarchical learning for complex manipulation skills with multisensory fusion," *Science Robotics*, vol. 4, no. 26, p. eaav3123, 2019.
- [13] T. S. Cohen and M. Welling, "Steerable CNNs," in *International Conference on Learning Representations*, 2017. [Online]. Available: <https://openreview.net/forum?id=rJQKYt5l1>
- [14] D. Wang and R. Walters, "So (2) equivariant reinforcement learning," in *International Conference on Learning Representations*, 2022.
- [15] D. Wang, R. Walters, X. Zhu, and R. Platt, "Equivariant q learning in spatial action spaces," in *Conference on Robot Learning*. PMLR, 2022, pp. 1713–1723.
- [16] M. Kalakrishnan, L. Righetti, P. Pastor, and S. Schaal, "Learning force control policies for compliant manipulation," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 4639–4644.
- [17] Y. Zhu, Z. Wang, J. Merel, A. Rusu, T. Erez, S. Cabi, S. Tunyasuvunakool, J. Kramár, R. Hadsell, N. de Freitas, and N. Heess, "Reinforcement and imitation learning for diverse visuomotor skills," in *Proceedings of Robotics: Science and Systems*, Pittsburgh, Pennsylvania, June 2018.
- [18] Y. Li, J.-Y. Zhu, R. Tedrake, and A. Torralba, "Connecting touch and vision via cross-modal prediction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10 609–10 618.
- [19] T. Cohen and M. Welling, "Group equivariant convolutional networks," in *International conference on machine learning*. PMLR, 2016, pp. 2990–2999.
- [20] M. Weiler and G. Cesa, "General e (2)-equivariant steerable cnns," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [21] G. Cesa, L. Lang, and M. Weiler, "A program to build e (n)-equivariant steerable cnns," in *International Conference on Learning Representations*, 2021.
- [22] T. S. Cohen, M. Geiger, J. Köhler, and M. Welling, "Spherical cnns," in *International Conference on Learning Representations*, 2018.
- [23] C. Kohler, N. Vaska, R. Muthukrishnan, W. Choi, J. Y. Park, J. Goodwin, R. Caceres, and R. Walters, "Symmetric models for radar response modeling," in *NeurIPS 2023 Workshop on Symmetry and Geometry in Neural Representations*, 2023.
- [24] N. Dym and H. Maron, "On the universality of rotation equivariant point cloud networks," in *International Conference on Learning Representations*, 2020.
- [25] D. Wang, J. Y. Park, N. Sortur, L. L. Wong, R. Walters, and R. Platt, "The surprising effectiveness of equivariant models in domains with latent symmetry," in *International Conference on Learning Representations*, 2023.
- [26] A. Simeonov, Y. Du, A. Tagliasacchi, J. B. Tenenbaum, A. Rodriguez, P. Agrawal, and V. Sitzmann, "Neural descriptor fields: Se (3)-equivariant object representations for manipulation," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 6394–6400.
- [27] X. Zhu, D. Wang, O. Biza, G. Su, R. Walters, and R. Platt, "Sample efficient grasp learning using equivariant models," *Proceedings of Robotics: Science and Systems (RSS)*, 2022.
- [28] H. Huang, D. Wang, R. Walters, and R. Platt, "Equivariant transporter network," in *Robotics: Science and Systems*, 2022.
- [29] S. Kim, B. Lim, Y. Lee, and F. C. Park, "Se (2)-equivariant pushing dynamics models for tabletop object manipulations," in *Conference on Robot Learning*. PMLR, 2023, pp. 427–436.
- [30] N. Dehmamy, R. Walters, Y. Liu, D. Wang, and R. Yu, "Automatic symmetry discovery with lie algebra convolutional network," *Advances in Neural Information Processing Systems*, vol. 34, pp. 2503–2515, 2021.
- [31] T. S. Cohen, M. Geiger, and M. Weiler, "A general theory of equivariant cnns on homogeneous spaces," *Advances in neural information processing systems*, vol. 32, 2019.
- [32] E. Van der Pol, D. Worrall, H. van Hoof, F. Oliehoek, and M. Welling, "Mdp homomorphic networks: Group symmetries in reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 4199–4210, 2020.
- [33] G. Cesa, L. Lang, and M. Weiler, "A program to build E(N)-equivariant steerable CNNs," in *International Conference on Learning Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=WE4qe9xlnQw>
- [34] M. Weiler and G. Cesa, "General E(2)-Equivariant Steerable CNNs," in *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [35] D. W. Romero and J.-B. Cordonnier, "Group equivariant stand-alone self-attention for vision," in *International Conference on Learning Representations*, 2020.
- [36] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *International conference on machine learning*. PMLR, 2018, pp. 1861–1870.
- [37] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double q-learning," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 30, no. 1, 2016.
- [38] D. Wang, C. Kohler, X. Zhu, M. Jia, and R. Platt, "Bulletarm: An open-source robotic manipulation benchmark and learning framework," in *The International Symposium of Robotics Research*. Springer, 2022, pp. 335–350.
- [39] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," 2016.
- [40] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," *arXiv preprint arXiv:1511.05952*, 2015.
- [41] A. X. Lee, A. Nagabandi, P. Abbeel, and S. Levine, "Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model," *Advances in Neural Information Processing Systems*, vol. 33, pp. 741–752, 2020.