Implicit Kinematic Policies: Unifying Joint and Cartesian Action Spaces in End-to-End Robot Learning

Aditya Ganapathi^{1,2}, Pete Florence¹, Jake Varley¹, Kaylee Burns^{1,3}, Ken Goldberg², Andy Zeng¹

Abstract—Action representation is an important yet often overlooked aspect in end-to-end robot learning with deep networks. Choosing one action space over another (e.g. target joint positions, or Cartesian end-effector poses) can result in surprisingly stark performance differences between various downstream tasks - and as a result, considerable research has been devoted to finding the right action space for a given application. However, in this work, we instead investigate how our models can discover and learn for themselves which action space to use. Leveraging recent work on implicit behavioral cloning, which takes both observations and actions as input, we demonstrate that it is possible to present the same action in multiple different spaces to the same policy - allowing it to learn inductive patterns from each space. Specifically, we study the benefits of combining Cartesian and joint action spaces in the context of learning manipulation skills. To this end, we present Implicit Kinematic Policies (IKP), which incorporates the kinematic chain as a differentiable module within the deep network. Experiments across several continuous control tasksfrom scooping piles of small objects, to lifting boxes with elbows, to precise block insertion with miscalibrated robots-suggest IKP not only learns complex prehensile and non-prehensile manipulation from pixels better than baseline alternatives, but also can learn to compensate for small joint encoder offset errors.

I. INTRODUCTION

Deep visuomotor policies that map from pixels to actions end-to-end can represent complex manipulation skills [15], but have shown to be sensitive to the choice of action space [17] - e.g., Cartesian end effector actions (i.e. task space [4]) perform favorably when learning policies for tabletop manipulation, while joint actions have shown to fare better for whole-body motion control [21], [28]. In particular, policies modeled by deep networks are subject to spectral biases [26], [3], [5] that make them more likely to learn and generalize the low-frequency patterns that exist in the control trajectories. Hence, choosing the right action space in which to define these trajectories remains a widely recognized problem in both reinforcement learning [21] as well as imitation learning, where demonstrations can be provided in a wide range of formats - e.g., continuous teleoperation in Cartesian or joint space, kinesthetic teaching, etc. - each changing the underlying characteristics of the training data.

Although considerable research has been devoted to finding the right action space for a given application [21], much less attention has been paid to figuring out how our models could instead discover for themselves which optimal combination of action spaces to use. This goal remains particularly unclear for most *explicit* policies $f_{\theta}(\mathbf{o}) = \hat{\mathbf{a}}$



Fig. 1. Enabled by an implicit policy formulation, Implicit Kinematic Policies (IKP) provide both Joint and Cartesian action representations (linked via forward kinematics) to a model that can pick up on inductive patterns in both spaces. Subsequently, this allows our end-to-end policies to extract the best action space for a wide variety of manipulation tasks ranging from tabletop block sorting to whole-arm sweeping.

with feed-forward models that take as input observations **o** and output actions $\mathbf{a} \in \mathcal{A}$. Explicit formulations must either specify a single action space to output from the model and potentially convert predictions into another desired space (e.g. with inverse kinematics), or use multiple action spaces with multiple outputs and losses which can be subject to conflicting gradients [33] from inconsistent predictions [17].

In this work, we demonstrate that we can train deep policies to learn which action space to use for imitation, made possible by implicit behavior cloning (BC) [8]. In contrast to its explicit counterpart, *implicit* BC formulates imitation as the minimization of an energy-based model (EBM) $\hat{\mathbf{a}} = \arg \min_{\mathbf{a} \in \mathcal{A}} E_{\theta}(\mathbf{o}, \mathbf{a})$ [14], which regresses the optimal action via sampling or gradient descent [32], [6]. Since actions are now inputs to the model, we propose presenting the model with the same action represented in multiple spaces while remaining consistent between each other, allowing the model to pick up on inductive patterns [3], [5] from all action representations.

We study the benefits of this multi-action-space formulation with Cartesian task and joint action spaces in the context of learning manipulation skills. Since both spaces are linked by the kinematic chain, we can integrate a differentiable kinematic module within the deep network – a formulation

¹Robotics at Google, ²UC Berkeley, ³Stanford University

which we refer to as Implicit Kinematic Policies (IKP). IKP can be weaved into an implicit autoregressive control strategy introduced in [8], where each action dimension is successively and uniformly sampled at a time and passed as input to the model. This exposes the model to not only the joint configuration and end-effector pose, but also the Cartesian action representations of *every rigid body link* of the arm as input. This can provide downstream benefits when learning whole-body manipulation tasks that may have emergent patterns in the trajectory of any given link of the robot. Furthermore, a key aspect of Implicit Kinematic Policies is that in addition to exposing the implicit policy to both Cartesian and joint action spaces, it also enables incorporating learnable layers throughout the kinematic chain, which can be used to optimize for residuals in either space.

Our main contribution is Implicit Kinematic Policies, a new formulation that integrates forward kinematics as a differentiable module within the deep network of an autoregressive implicit policy, exposing both joint and Cartesian action spaces in a kinematically consistent manner. Behavior cloning experiments across several complex visionbased manipulation tasks suggest that IKP, without any modifications, is broadly capable of efficiently learning both prehensile and non-prehensile manipulation tasks, even in the presence of miscalibrated joint encoders - results that may pave the way for more data efficient learning on low-cost or cable-driven robots, where low-level joint encoder errors (due to drift, miscalibration, or gear backlash) can propagate to large non-linear artifacts in the Cartesian end effector trajectories. IKP achieves 85.9%, 97.5% and 92.4% on a sweeping, non-prehensile box lifting and precise insertion task respectively - performance that is on par with or exceeds that of standard implicit or explicit baselines (using the best empirically chosen action space per individual task). We also provide qualitative experiments on a real UR5e robot demonstrating IKP's ability to learn from noisy human demonstrations in a data-efficient manner. Our experiments in both simulated and real environments provide an initial case study towards more general action-space-agnostic architectures for end-to-end robot learning. Code will be made available.

II. RELATED WORK

A. Learning Task-Specific Action Representations

While much of the early work in end-to-end robot learning has focused on learning explicit policies that map pixels to low-level joint torques [16], [15], more recent work has found that the choice of action space can have a significant impact on downstream performance for a wide range of robotic tasks ranging from manipulation to locomotion [18], [7], [21], [31], [30]. In the context of deep reinforcement learning for locomotion, [21] find that using PD joint controllers achieves higher sample efficiency and reward compared to torque controllers on bipedal and quadrupedal walking tasks in simulation, whereas [31] interestingly find that torque controllers outperform PD joint controllers for a real robot hopping task. In contrast, [7] shows that using task space actions are more effective for learning higher level goals of legged locomotion compared to joint space controllers and demonstrate results on a real bipedal robot.

Careful design of the action space is equally important when learning manipulation skills with high DoF robots [18], [30]. Prior approaches successfully use direct torque control or PD control with joint targets to learn a variety of tasks such as pick-and-place, hammering, door opening [25] and stacking [23], but these methods often require lots of data and interaction with the environment to generalize well and are also sensitive to feedback gains in the case of PD control.

Other works default to low-level residual cartesian PD controllers in order to learn prehensile tasks such as block stacking [13] and insertion [27], but [18], [30] instead vouch for a task-space impedance controller to support more compliant robot control. However, this approach is agnostic to the full joint configuration of the robot and focuses on tabletop manipulation which is generally better defined in task space. Fundamentally, these works share the common theme of choosing the single best action space for a particular task. In this work, we are more concerned with the problem of extracting the best combination of action spaces.

B. Difficulties in Learning High-Frequency Functions

The proposed benefits of implicit kinematic policies rely on the hypothesis that simpler patterns are easier for deep networks to learn. Although theoretical analyses of two-layer networks and empirical studies of random training labels have shown that deep networks are prone to memorization [34], further studies on network convergence behavior and realistic dataset noise have found that this result does not necessarily explain performance in practice. Instead, deep networks first learn patterns across samples before memorizing data points [1]. [24] formalize this finding and show that deep networks have a learning bias towards low-frequency functions by using tools from Fourier analysis and [2] extend this result by showing that networks fit increasingly higher frequency functions over the course of training. Within learning-based robotics, multiple works have found the related result that higher-level action spaces lead to improved learnability. Specifically, the frequency of the action space can have an outsized effect on the robustness and quality of learned policies [21]. This could explain why learned, latent actions improve performance and generalizability for policies trained either with reinforcement learning [10], [11], [19], [22] or imitation learning [9].

III. BACKGROUND

A. Problem Formulation

We consider the imitation learning setting with access to a fixed dataset of observation-action pairs. We frame this problem as supervised learning of a policy $\pi : \mathcal{O} \to \mathcal{A}$ where $\mathbf{o} \in \mathcal{O} = \mathbb{R}^m$ represents an observational input to the policy and $\mathbf{a} \in \mathcal{A} = \mathbb{R}^d$ represents the action output. The policy π , with parameters θ , can either be formulated as an *explicit* function $f_{\theta}(\mathbf{o}) \to \hat{\mathbf{a}}$ or as an *implicit* function $\arg \min_{\mathbf{a} \in \mathcal{A}} E(\mathbf{o}, \mathbf{a}) \to \hat{\mathbf{a}}$. In this work, we adopt the



Fig. 2. Integrating forward kinematics (FK), in particular, synergizes with implicit policies (a, b), where actions **a** are sampled in joint space, and both joints and corresponding Cartesian actions along with observations **o** can be fed as inputs to the EBM f_{θ} . For explicit policies (c), it becomes less clear how to expose both action spaces in a kinematically consistent way to the deep network, except via predicting joints from observations, then using forward kinematics to backpropagate gradients from losses imposed on the final output Cartesian actions (in red). In this work, we investigate both formulations and find that implicit (b) yields better performance.

implicit formulation and build upon recent work which we describe in more detail in Section III-B. We operate in the continuous control setting, where our policy infers a desired target configuration at 10 Hz, and a low-level controller asynchronously reaches desired configurations with a joint-level PD controller. Our \mathcal{O} contains both image observations and proprioceptive robot state (i.e. from joint encoders).

B. Implicit BC

We use an energy-based, contrastive loss to train our implicit policy, modeled with the same late-fusion deep network architecture as in [8] with 26 convolutional ResNet layers [12] for the image encoder and 20 dense ResNet layers for the EBM. Specifically, to optimize $E_{\theta}(\cdot)$, we train an InfoNCE style loss [29] on observation-action samples [8]. Our dataset consists of $\{\mathbf{o}_i, \mathbf{a}_i^*\}$ for $\mathbf{o}_i \in \mathbb{R}^m$ and $\mathbf{a}_i \in \mathbb{R}^d$ and from regression bounds $\mathbf{a}_{\min}, \mathbf{a}_{\max} \in \mathbb{R}^d$ we generate negative samples $\{\tilde{\mathbf{a}}_i^j\}_{j=1}^{N_{neg.}}$. We estimate $p(\mathbf{a}|\mathbf{o})$ by using our negative samples to approximate our normalizing constant:

$$\begin{split} \mathcal{L}_{\text{InfoNCE}} &= \sum_{i=1}^{N} -\log(\tilde{p}_{\theta}(\mathbf{a}_{i}^{*}|\mathbf{o}_{i}, \{\tilde{\mathbf{a}}_{i}^{j}\}_{j=1}^{N_{\text{neg.}}}))\\ \tilde{p}_{\theta}(\mathbf{a}_{i}^{*}|\mathbf{o}_{i}, \{\tilde{\mathbf{a}}_{i}^{j}\}_{j=1}^{N_{\text{neg.}}}) &= \frac{e^{-E_{\theta}(\mathbf{o}_{i},\mathbf{a}_{i}^{*})}}{e^{-E_{\theta}(\mathbf{o}_{i},\mathbf{a}_{i}^{*})} + \sum_{j=1}^{N_{\text{neg.}}} e^{-E_{\theta}(\mathbf{o}_{i},\tilde{\mathbf{a}}_{i}^{j})}} \end{split}$$

To perform inference, given some observation \mathbf{o} , we minimize our learned energy function, $E_{\theta}(\mathbf{o}, \mathbf{a})$, over actions in \mathcal{A} using a sampling-based optimization procedure.

IV. METHOD

Here we present the details of our proposed Implicit Kinematic Policies (IKP) method. First, we will present our extension to the implicit policy formulation which provides multiple action space representations as input. (Sec. IV-A). Sec. IV-A.1 also compares the implicit multi-actionspace formulation to its explicit counterpart and discusses the relevant tradeoffs. We then describe how to perform autoregressive training and inference with implicit policies in a way that exposes the action trajectories of every joint and link in the robot to the model (Sec. IV-B). Finally, we discuss a motivating application of controlling miscalibrated robots for precise tasks, and how IKP is uniquely suited to automatically compensate for this noise through the use of strategically placed residual blocks in the model (Sec. IV-C).

A. Multi-Action-Spaces With Implicit Policies

Our formulation enables the implicit policy $\hat{\mathbf{a}} = \arg \min_{\mathbf{a}} E_{\theta}(\mathbf{o}, \mathbf{a})$ to have access to multiple action spaces, which are constrained to be consistent. Specifically, our implicit multi-action-space formulation is of the form:

$$\underset{\mathbf{a},\mathbf{a}',\dots}{\operatorname{arg\,min}} \quad E_{\theta}(\mathbf{o},\mathbf{a},\mathbf{a}',\dots)$$

$$\underset{\mathbf{s},\mathbf{t}, \mathbf{a}}{\operatorname{st}} = \mathcal{T}(\mathbf{a}'),\dots$$

$$(1)$$

where $\mathbf{a} \in \mathcal{A}$ is one parameterization of the action space, $\mathbf{a}' \in \mathcal{A}'$ is a different parameterization of the action space, and $\mathcal{T}(\cdot) : \mathcal{A} \to \mathcal{A}'$ is a transformation between the two action spaces. Of course, N different consistent parameterizations of the action space could be represented, which is depicted in Eq. 1 by the ellipses (...).

In particular for robots, we are interested in representing both *joint-space* and *cartesian-space* actions, which is a case of Eq. 1 in which the transformation between the two action spaces is forward kinematics (FK):

а

$$\begin{array}{ll} \arg\min_{\text{joints, } \mathbf{a}_{\text{cartesian}}} & E_{\theta}(\mathbf{o}, \mathbf{a}_{\text{joints}}, \mathbf{a}_{\text{cartesian}}) \\ \text{s.t.} & \mathbf{a}_{\text{cartesian}} = FK(\mathbf{a}_{\text{joints}}) \end{array}$$
(2)

Specifically, we accomplish this by first sampling input actions in joint space $\mathbf{a}_{\text{joints}} \in \mathcal{A}_{\text{joints}}$, then computing the corresponding Cartesian (task) space actions via forward kinematics (*FK*), $\mathbf{a}_{\text{cartesian}} = FK(\mathbf{a}_{\text{joints}})$, and finally concatenating and passing the combined representation into the model $E_{\theta}(\mathbf{o}, \mathbf{a}_{\text{joints}}, \mathbf{a}_{\text{cartesian}})$. A visualization of this model is shown in the middle portion of Fig. 2. Our hypothesis is that the model can use this redundant action representation to exploit patterns in both spaces, akin to *automatically* discovering the best combination of action spaces – this hypothesis will be tested in our Experiments section. As we will show in Sec. IV-B, we can perform training and inference for E_{θ} through using autoregressive derivative-free optimization.

1) Multi-Action-Spaces With Explicit Policies: Consider the above formulation in contrast to an explicit policy, $\hat{\mathbf{a}} = f_{\theta}(\mathbf{o})$, where the mapping $f_{\theta}(\cdot)$ maps to one specific action space. It is possible to provide different losses on different transformations of the explicit policy's action space, for example $\mathcal{L}_{\text{joints}}(\hat{\mathbf{a}}, \mathbf{a}^*)$ and $\mathcal{L}_{\text{cartesian}}(FK(\hat{\mathbf{a}}), FK(\mathbf{a}^*))$, where FK represents differentiable forward kinematics. However, this can lead to conflicting gradients and requires choosing relative weightings, λ , between these losses: $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{joints}} + \lambda \mathcal{L}_{\text{cartesian}}$. This issue of loss balancing is exacerbated if the explicit policy were to regress the Cartesian representation of every link in the robot in addition to the joint configuration and end-effector pose. Additionally, the explicit policy predicts a single joint configuration even though there may be multiple correct configurations when using kinematically



Fig. 3. **Method overview.** Given an image captured from an RGB camera overlooking the robot workspace (a), we feed it as input to a deep convolutional network (b) to get a latent state representation. We then predict desired robot joint actions by leveraging implicit autoregression [8], [20] (c) with state-conditioned EBMs to progressively sample each action dimension (joint angle) at a time: i.e., we uniform sample j_n , feed it to FK_n (forward kinematics to link *n*, prepended with deep layers) to get the Cartesian pose C_n , which is then concatenated with the latent state and all previously sampled argmin joint dimensions $[j_n, j_{n-1}, ..., j_0]$ and Cartesian representations $[C_{n-1}, ..., C_0]$ and fed to an 8-layer EBM E_n to compute the argmin over j_n .

redundant robots. It is possible to first regress Cartesian actions and subsequently recover the corresponding joint configuration via inverse kinematics, however, IK solvers are only approximately differentiable which poses challenges for end-to-end training. A visualization of this explicit policy is shown in the rightmost column of Fig. 2.

B. Autoregressive Implicit Training and Inference

As in [8], we employ an autoregressive derivativefree optimization method for both training and inference. Sampling-based, derivative-free optimization synergizes with the multi-action-space constrained energy model (Eq. 1), since gradient-based optimization (i.e. through Langevin dynamics [8]), while possible, would require a solution to synchronize between the different action spaces.²

The autoregressive procedure uses m models, one model $E^j_{\theta}(\mathbf{o}, \mathbf{a}^{0:j})$ for each dimension j = 1, 2, ..., m. In contrast with prior work [8], instead of \mathbf{a} only representing a single action space, our proposed method represents, for each *j*th model, both the 6DoF Cartesian pose of the *j*th link *and* the joint of that link into \mathbf{a} . Similarly in \mathbf{o} we also represent the dual Cartesian-and-joint state of each link. Since our method strictly samples actions in joint space and the state is represented by the joint configuration, the Cartesian representations for both the state \mathbf{o} and actions \mathbf{a} are acquired through $FK(\mathbf{o})$ and $FK(\mathbf{a})$ for every dimension j = 1, 2, ..., m. We visualize this procedure in Fig. 3.

C. Residual Forward Kinematics

In robotics, we often assume access to accurate kinematic descriptions and joint encoders, but this assumption can be a significant source of error in the context of low-cost or cable-driven robots. In these settings, the errors often take the form of fixed, unknown linear offsets in each joint motor encoder or link representation. If calibration isn't performed, both cases can cause heavy non-linear offsets in end-effector space which can drastically affect the performance on highprecision tasks. By adding dense residual blocks both before and after joint actions are sampled in the autoregressive procedure, we can encourage our model to learn these linear offsets at each joint and link rather than solving the more difficult problem of learning a highly non-linear offset in task space. This extension is only possible due to the full differentiability of the forward kinematics layer, which allows gradients to flow through the residual blocks. Concretely, we redefine our multi-action-space formulation with residuals as follows:

$$\begin{array}{ll} \underset{\mathbf{a}_{\text{joints}}, \mathbf{a}_{\text{cartesian}}}{\arg\min} & E_{\theta}(\mathbf{o}, \mathbf{a}_{\text{joints}}, \mathbf{a}_{\text{cartesian}}) \\ \text{s.t.} & \mathbf{a}_{\text{cartesian}} = FK(\mathbf{a}_{\text{joints}} + \Delta_{\theta}(\mathbf{a}_{\text{joints}})) \end{array}$$
(3)

where $\Delta_{\theta}(\mathbf{a}_{joints}) \in \mathcal{A}_{joints}$ is a new learnable module in the implicit model which gives the model the inductive bias that there may be imperfect calibration of the joint measurements (i.e., biased joint encoder measurements). The module $\Delta_{\theta}(\cdot)$ can be learned during training and we hypothesize that the EBM may be able to automatically learn the $\Delta_{\theta}(\cdot)$ which provides the lowest-energy fit of the data. This can be interpreted as automatically calibrating the biased joint encoders. Experiments testing this hypothesis are in Section V-B.

V. EXPERIMENTS

We evaluate IKP across several vision-based continuous control manipulation tasks with quantitative experiments against baseline methods in simulation, as well as qualitative results on a real robot. All tasks require generalization to unseen object configurations at test time. The goals of our experiments are two-fold: (i) across tasks where one action space substantially outperforms the other, we investigate whether IKP can achieve the best of both and perform consistently well across all tasks, and (ii) given a miscalibrated robot with unknown offsets in the joint encoders (representing miscalibration or low-cost encoders), we study whether IKP can autonomously learn to compensate for these offsets while still succeeding at the task. Across all experiments, a dataset of expert demonstrations is provided by a scripted oracle policy (in simulation), or by human teleoperation (in real).

²Specifically, when using the implicit model gradients, with learning rate λ , the actions are updated as $\mathbf{a}_{\text{joints}}^{k+1} = \mathbf{a}_{\text{joints}}^k - \lambda \nabla_{\mathbf{a}_{\text{joints}}} E_{\theta}(\cdot)$ and $\mathbf{a}_{\text{cartesian}}^{k+1} = \mathbf{a}_{\text{cartesian}}^k - \lambda \nabla_{\mathbf{a}_{\text{cartesian}}} E_{\theta}(\cdot)$. The updated actions actions may not be consistent with each other – even for an analytical model, there will be differences in the first-order derivatives of the different spaces.



Fig. 4. Expert trajectories for bimanual sweeping are characterized by distinct patterns in Cartesian space (b), e.g., y-values experience a modechange in the latter half of the episode when the policy switches bowls in which to drop particles. Such patterns are less salient in joint space (a). The opposite is true for bimanual flipping, where linear patterns emerge in joint space (c), whereas they are less salient in Cartesian space (d) especially with randomized end effector poses (semi-transparent plots).

A. Simulated Bimanual Sweeping and Flipping

In a simulated environment, we evaluate on two bimanual tasks - sweeping and flipping - both of which involve two 7DoF KUKA IIWA robot arms equipped with spatulalike end-effectors positioned over a $0.4m^2$ workspace. The setup for sweeping is identical to that presented in Florence et al. [8], where given a pile of 50-60 particles on the workspace, the task is to scoop and evenly distribute all particles into two bowls located next to the workspace. For flipping, given a bowl of 20-30 particles attached to the top of a $0.2m^3$ box, the task is to flip the box to pour the particles into a larger bowl positioned near the base of the arms. Both tasks are designed to involve significant coordination between the two arms. During sweeping, for example, scooping up particles and transporting them to the bowls requires carefully maintaining alignment between the tips of the spatulas to avoid dropping particles. Flipping, on the other hand, requires aligning the elbow surfaces against the sides of the box, and using friction to carry the box as it pours the particles into the bowl - any subtle misalignment of the elbow against the sides can lead to the box slipping away or tipping over. For both tasks, we generate fixed datasets of 1,000 demonstration episodes for imitation learning using a scripted oracle with access to privileged state information, including object poses and contact points, which are not accessible to policies.

We train two Implicit BC [8] baselines *without* Residual Kinematics: one using a 12DoF Cartesian action space (6DoF end-effector pose for each arm), and another using a 14DoF joint action space (7DoF for each arm). From the quantitative results presented in Tab. I, we observe that the performance of Implicit BC on bimanual sweeping degrades substantially when using a joint action space. This is likely that learning the task (i.e. distribution of oracle demonstrations) involves recognizing a number of Cartesian space constraints that are less salient in joint space: e.g., keeping the Cartesian poses of the spatulas aligned with each other while scooping and transporting particles, and maintaining z-height with

Fig. 5. Simple constant joint encoder errors (which may appear from drift in low-cost or cable-driven robots) can propagate to non-linear offsets in Cartesian space. By learning joint space residuals to compensate for these offsets, IKP is better equipped to generalize over such errors than standard end-to-end policies trained in Cartesian space for tabletop manipulation.

the tabletop. For bimanual flipping, on the other hand, we observe the opposite: where the performance of Implicit BC performs well with joint actions, but poorly with Cartesian actions. We conjecture that the task involves more wholebody manipulation, where fitting policies in joint space are more likely to result in motion trajectories that accurately conform to the desired elbow contact with the box.

| TABLE I | |
|---------|--|
| | |

Performance Measured in Task Success (AVG \pm STD % Over 3 Seeds)

| Task Oracle's action space Policy Action Space | Sweeping cartesian | Flipping joints |
|--|---|---|
| Cartesian Joints Joints + Cartesian (Ours) | $\begin{array}{c} 79.4 \pm 2.1 \\ 44.3 \pm 3.2 \\ \textbf{85.9} \pm \textbf{1.5} \end{array}$ | $\begin{array}{c} 38.6 \pm 5.2 \\ \textbf{98.4} \pm \textbf{1.4} \\ \textbf{97.5} \pm \textbf{1.2} \end{array}$ |

Our proposed method, Implicit Kinematic Policies (labeled as Joints + Cartesian in Tab. I) leverages both action spaces. Results suggest its performance is not only on par with the best bespoke action space for the task, but also surprisingly exceeds the performance of Cartesian actions for bimanual sweeping. Upon further inspection of the action trajectories for bimanual sweeping, it is evident that the Cartesian space action trajectories (Fig. 4b) are lower frequency and contain more piecewise linear structure in comparison to the joint space action trajectories (Fig. 4a). Implicit policies have shown to thrive in conditions where this structure is present [8]. In contrast, for bimanual flipping, the joint space action trajectories (Fig. 4c) appear much more linear and low frequency than the corresponding arc-like Cartesian trajectories (Fig. 4d). We also add small perturbations to the end-effector pose highlighted by the semi-transparent lines in plot (d) of Fig. 4. Note that the Implicit (Cartesian) policy uses a generic IK solver which can return multiple distinct joint target commands, causing less stable trajectories overall.

B. Simulated Miscalibrated Sweeping and Insertion

We design two additional tasks to evaluate IKP's ability to learn high precision tasks in the presence of inaccurate joint encoders, described in Sec. IV-C. To replicate such a scenario, we simulate these errors by adding constant < 2 degree offsets to each of the six revolute joints on a UR5e robot. These offsets are visualized and described in more detail in Fig. 5. We test the effect of these offsets on the bimanual sweeping task and a new insertion task where the goal is to insert an L-shaped block into a tight fixture. Demonstrations for both tasks are provided in the form of a Cartesian scripted oracle with privileged access to the underlying joint offsets and can compensate for them. This is akin to a human using visual servoing to compensate for inaccurate encoders when teleoperating a real robot to collect expert demonstrations. Although both tasks are generally well-defined in Cartesian space as a series of linear step functions, the joint offsets induce high-frequency non-linear artifacts in the end-effector trajectory (Fig. 5b and 5d) causing poor performance with the Implicit (Cartesian) policy (results in Tab. II). Poor performance persists for miscalibrated insertion despite using 10x the data. Alternatively, even though the encoder offsets only cause linear shifts in the joint trajectories (Fig. 5a and 5c), both tasks are less structured in joint space and as a result Implicit (Joints) performs significantly worse on miscalibrated bimanual sweeping and comparatively worse on miscalibrated insertion. Implicit (Joints) performs relatively well on the miscalibrated insertion task due to the simplicity of the pick and place motion, but struggles to generalize when the block nears the edge of workspace as the joint trajectories become increasingly non-linear in those regions. The performance only slightly improves when using 1000 demonstrations. IKP provides a best-of-all-worlds solution by learning the linear offsets in joint space through the residual blocks while simultaneously exploiting the unperturbed Cartesian trajectories through forward kinematics on the shifted joint actions to generalize. IKP achieves the highest performance on miscalibrated insertion with both 100 and 1000 demonstrations, and significantly higher performance on miscalibrated bimanual sweeping. Interestingly, not only do Explicit and ExplicitFK perform significantly worse than both Implicit (Joints) and IKP for both tasks, but ExplicitFK also provides little to no additional benefit over Explicit, even in the presence of more data.

TABLE II

Miscalibrated Joint Encoder Experiments, Performance Measured in Task Success (Avg \pm Std. % Over 3 Seeds). ("J+C") is short for

JOINTS+CARTESIAN.

| Task | Sweeping | Block Insertion | |
|-----------------------------|----------------|----------------------------------|----------------|
| Oracle's action space | cartesian | cartesian | |
| Method (Action Space) | 1000 | 100 | 1000 |
| Explicit (Joints) | 38.2 ± 3.4 | 73.8 ± 4.3 | 74.2 ± 2.5 |
| ExplicitFK $(J + C)$ (Ours) | 40.3 ± 4.6 | 72.4 ± 3.8 | 75.4 ± 2.8 |
| Implicit (Cartesian) | 3.1 ± 2.2 | 0.0 ± 0.0 | 0.0 ± 0.0 |
| Implicit (Joints) | 46.3 ± 1.8 | 82.3 ± 2.9 | 85.2 ± 3.1 |
| IKP (J + C) (Ours) | 84.5 ± 1.2 | $\textbf{88.8} \pm \textbf{3.4}$ | 92.4 ± 2.6 |

C. Real Robot Sorting, Sweeping, and Alignment

We conduct qualitative experiments with a real UR5e robot on two tasks: 1) sorting two blocks into bowls and



Fig. 6. Real robot IKP rollouts on a sweeping and sorting task (top two rows) and a block alignment task (bottom row).

subsequently sweeping the bowls, and 2) aligning a red block with a blue block, where both tasks have initial locations randomized). Our goal with these experiments is two-fold: i) to demonstrate that IKP can run on real robots with noisy human continuous teleop demonstrations with as few as 100 demonstrations, and ii) to show that Implicit Kinematic Policies can perform both whole-body and prehensile tasks. Demonstrations for the block sorting portion of the first task are provided using continuous Cartesian position control based teleoperation at 100 Hz with a 3D mouse. Once the blocks are sorted, the teleoperator switches to a joint space PD controller with the same 3D mouse to guide the arm into an extended pose in order to sweep the bowls. For the alignment task, demonstrations are also provided using x, y and z Cartesian PD control at 100 Hz. Both tasks take 480×640 RGB images (downsampled to 96×96) from an Intel RealSense D435 camera at a semi-overhead view as input to the policy, and no extrinsic camera calibration is used. We show successful rollouts for both tasks in Fig. 6. We plan to further investigate quantitative experiments after COVID-related restrictions relax.

VI. DISCUSSION

In future work, we will investigate using the multi-actionspace formulation to extend beyond joint and cartesian PD control by incorporating the underlying forward and/or inverse robot dynamics into the network which will allow us to expose joint torques and velocities to our model in a fully differentiable way. We hypothesize that this additional information will be particularly helpful in dynamic manipulation tasks (e.g. flipping a pancake) and visual locomotion where torque and velocity action trajectories may appear more structured. We would also like to further test our residual framework within IKP on robots that have nonlinear drift in joint space which may represent a larger set of low-cost robots. Finally, we would like to utilize our fully differentiable residual forward kinematics module to learn the link parameters themselves, which can be a promising direction for controlling soft and/or continuum robots.

ACKNOWLEDGMENTS

The authors would like to thank Vincent Vanhoucke, Vikas Sindhwani, Johnny Lee, Adrian Wong, Daniel Seita and Ryan Hoque for helpful discussions and valuable feedback on the manuscript.

REFERENCES

- D. Arpit, S. Jastrzebski, N. Ballas, D. Krueger, E. Bengio, M. S. Kanwal, T. Maharaj, A. Fischer, A. C. Courville, Y. Bengio, and S. Lacoste-Julien, "A closer look at memorization in deep networks," *ICML*, vol. abs/1706.05394, 2017.
- [2] R. Basri, D. Jacobs, Y. Kasten, and S. Kritchman, "The convergence rate of neural networks for learned functions of different frequencies," in *NeurIPS*, 2019.
- [3] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner *et al.*, "Relational inductive biases, deep learning, and graph networks," *arXiv preprint arXiv:1806.01261*, 2018.
- [4] D. Berenson, S. Srinivasa, and J. Kuffner, "Task space regions: A framework for pose-constrained manipulation planning," *The International Journal of Robotics Research*, vol. 30, no. 12, pp. 1435–1460, 2011.
- [5] A. Bietti and J. Mairal, "On the inductive bias of neural tangent kernels," *arXiv preprint arXiv:1905.12173*, 2019.
- [6] Y. Du and I. Mordatch, "Implicit generation and modeling with energy based models," 2019.
- [7] H. Duan, J. Dao, K. Green, T. Apgar, A. Fern, and J. Hurst, "Learning task space actions for bipedal locomotion," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [8] P. Florence, C. Lynch, A. Zeng, O. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson, "Implicit behavioral cloning," arXiv preprint arXiv:2109.00137, 2021.
- [9] R. Fox, S. Krishnan, I. Stoica, and K. Goldberg, "DDCO: Discovery of Deep Continuous Options for Robot Learning from Demonstrations," in *Conference on Robot Learning (CoRL)*, 2017.
- [10] T. Haarnoja, K. Hartikainen, P. Abbeel, and S. Levine, "Latent space policies for hierarchical reinforcement learning," in *ICML*, 2018.
- [11] K. Hausman, J. T. Springenberg, Z. Wang, N. Heess, and M. A. Riedmiller, "Learning an embedding space for transferable robot skills," in *ICLR*, 2018.
- [12] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [13] T. Johannink, S. Bahl, A. Nair, J. Luo, A. Kumar, M. Loskyll, J. A. Ojea, E. Solowjow, and S. Levine, "Residual reinforcement learning for robot control," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2019.
- [14] Y. LeCun, S. Chopra, R. Hadsell, M. Ranzato, and F. Huang, "A tutorial on energy-based learning," *Predicting structured data*, vol. 1, no. 0, 2006.
- [15] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end Training of Deep Visuomotor Policies," in *Journal of Machine Learning Research* (*JMLR*), 2016.
- [16] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous Control with Deep Reinforcement Learning," in *International Conference on Learning Representations* (*ICLR*), 2016.
- [17] R. Martín-Martín, M. Lee, R. Gardner, S. Savarese, J. Bohg, and A. Garg, "Variable impedance control in end-effector space. an action space for reinforcement learning in contact rich tasks," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [18] R. Martín-Martín, M. A. Lee, R. Gardner, S. Savarese, J. Bohg, and A. Garg, "Variable impedance control in end-effector space: An action space for reinforcement learning in contact-rich tasks," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [19] S. Nair and C. Finn, "Hierarchical foresight: Self-supervised learning of long-horizon tasks via visual subgoal generation," *ArXiv*, vol. abs/1909.05829, 2020.
- [20] C. Nash and C. Durkan, "Autoregressive energy machines," in *Inter*national Conference on Machine Learning. PMLR, 2019, pp. 1735– 1744.
- [21] X. B. Peng and M. van de Panne, "Learning locomotion skills using deeptl: Does the choice of action space matter?" in ACM SIGGRAPH Eurographics Symposium on Computer Animation (SCA), 2017.
- [22] M. Pflueger and G. Sukhatme, "Plan-space state embeddings for improved reinforcement learning," ArXiv, vol. abs/2004.14567, 2020.
- [23] I. Popov, N. Heess, T. Lillicrap, R. Hafner, G. Barth-Maron, M. Vecerik, T. Lampe, Y. Tassa, T. Erez, and M. Riedmiller, "Data-efficient deep reinforcement learning for dexterous manipulation," 2018.

- [24] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio, and A. Courville, "On the spectral bias of neural networks," in *Proceedings of the 36th International Conference* on Machine Learning, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 09–15 Jun 2019, pp. 5301–5310. [Online]. Available: https://proceedings.mlr.press/v97/rahaman19a.html
- [25] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," in *Robotics: Science and Systems (RSS)*, 2018.
- [26] B. Ronen, D. Jacobs, Y. Kasten, and S. Kritchman, "The convergence rate of neural networks for learned functions of different frequencies," *Advances in Neural Information Processing Systems*, vol. 32, pp. 4761–4771, 2019.
- [27] T. Silver, K. Allen, J. Tenenbaum, and L. Kaelbling, "Residual policy learning," 2019.
- [28] J. Tan, T. Zhang, E. Coumans, A. Iscen, Y. Bai, D. Hafner, S. Bohez, and V. Vanhoucke, "Sim-to-real: Learning agile locomotion for quadruped robots," in *Proceedings of Robotics: Science and Systems*, Pittsburgh, Pennsylvania, June 2018.
- [29] A. van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *ArXiv*, vol. abs/1807.03748, 2018.
- [30] P. Varin, L. Grossman, and S. Kuindersma, "A comparison of action spaces for learning manipulation tasks," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [31] J. Viereck, J. Kozolinsky, A. Herzog, and L. Righetti, "Learning a structured neural network policy for a hopping task," vol. 3, no. 4. Institute of Electrical and Electronics Engineers (IEEE), Oct 2018, p. 4092–4099. [Online]. Available: http://dx.doi.org/10.1109/LRA.2018.2861466
- [32] M. Welling and Y. W. Teh, "Bayesian learning via stochastic gradient langevin dynamics," in *Proceedings of the 28th international conference on machine learning (ICML-11)*. Citeseer, 2011, pp. 681–688.
- [33] T. Yu, S. Kumar, A. Gupta, S. Levine, K. Hausman, and C. Finn, "Gradient surgery for multi-task learning," in *Neural Information Processing Systems (NeurIPS)*, 2020.
- [34] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, "Understanding deep learning requires rethinking generalization," *ICLR*, vol. abs/1611.03530, 2017.