

Appendices

395 A KODex Pseudo-code

The overall pseudo-code for KODex is shown below.

Algorithm 1: KODex

Demonstration Data Collection

Initialize $D = \emptyset$;

for $n \in \{1, \dots, N\}$ **do**

Generate a $T^{(n)}$ -horizon trajectory of states and torques $\{[x^n(t), \tau^n(t)]\}_{t=1}^{t=T^{(n)}}$;

Add $\{[x^n(t), \tau^n(t)]\}_{t=1}^{t=T^{(n)}}$ to D ;

end

Koopman Operator Approximation

Determine lifting function $\phi(x(t))$;

Compute \mathbf{K} on D (6, 9);

Controller Design

Build a controller C as a neural network with inputs as $(x_r(t), x_r(t+1))$ and output as $\tau(t)$;

Train C using state-torque pairs $(x_r^n(t), x_r^n(t+1), \tau^n(t))$ in D (10);

Implementation

Specify the initial states $x(1)$;

for $t \in \{1, \dots, T-1\}$ **do**

Predict the next robot states $\hat{x}_r(t+1)$ using \mathbf{K} (3 8);

Read the current robot states $x_r(t)$;

Generate the torque $\tau(t)$ using C on $(x_r(t), \hat{x}_r(t+1))$ and execute it;

end

397 B State Design

398 In this section, we show the state design for each task in detail. It should be noted that the motion
 399 capability of the hand for each task were suggested from the work [8] that originally introduced
 400 these tasks. For a decent implementation, we employed the same setting.

401 **Tool use** For this task, the floating wrist base can only rotate along the x and y axis, so we have
 402 $x_r(t) \in \mathcal{X}_r \subset \mathbb{R}^{26}$. Regarding the object states, unlike the other tasks, where the objects of
 403 interest are directly manipulated by the hand, this task requires to modify the environment it-
 404 self. As a result, except for the hammer positions, orientations and their corresponding veloci-
 405 ties $p_t^{\text{tool}}, o_t^{\text{tool}}, \dot{p}_t^{\text{tool}}, \dot{o}_t^{\text{tool}}$ (\mathbb{R}^3), we also define the nail goal position p^{nail} (\mathbb{R}^3). Finally, we have
 406 $x_o(t) = [p_t^{\text{tool}}, o_t^{\text{tool}}, \dot{p}_t^{\text{tool}}, \dot{o}_t^{\text{tool}}, p^{\text{nail}}] \in \mathcal{X}_o \subset \mathbb{R}^{15}$. As a result, $x(t)$ includes 41 states in total and we
 407 use $T = 100$.

408 **Door opening** For this task, the floating wrist base can only move along the direction that is perpen-
 409 dicular to the door plane but rotate freely, so we have $x_r(t) \in \mathcal{X}_r \subset \mathbb{R}^{28}$. Regarding the object states,
 410 we define the fixed door position p^{door} , which can provide with case-specific information (similar to
 411 p^{nail} in Tool Use), and the handle positions p_t^{handle} (both \mathbb{R}^3). In order to take into consideration the
 412 status of door being opened, we include the angular velocity of the opening angle v_t (\mathbb{R}^1). Finally,
 413 we have $x_o(t) = [p_t^{\text{handle}}, v_t, p^{\text{door}}] \in \mathcal{X}_o \subset \mathbb{R}^7$. As a result, $x(t)$ includes 35 states in total and we
 414 use $T = 70$.

415 **Object relocation** For this task, the ADROIT hand is fully actuated, so we have $x_r(t) \in \mathcal{X}^r \subset \mathbb{R}^{30}$
 416 (24-DoF hand + 6-DoF floating wrist base). Regarding the object states, we define p^{target} and p_t^{ball}
 417 as the target and current positions. Then, we compute $\bar{p}_t^{\text{ball}} = p_t^{\text{ball}} - p^{\text{target}}$, which is the component
 418 of p_t^{ball} in a new coordinate frame that is constructed by p^{target} being the origin. We additional in-
 419 clude the ball orientation o_t^{ball} and their corresponding velocities $\dot{p}_t^{\text{ball}}, \dot{o}_t^{\text{ball}}$ (all \mathbb{R}^3). Finally, we have
 420 $x_o(t) = [\bar{p}_t^{\text{ball}}, o_t^{\text{ball}}, \dot{p}_t^{\text{ball}}, \dot{o}_t^{\text{ball}}] \in \mathcal{X}_o \subset \mathbb{R}^{12}$. As a result, $x(t)$ includes 42 states in total and we use

421 $T = 100$.

422 **In-hand reorientation** For this task, the floating wrist base is fixed, so we only consider the 24-DoF
423 hand joints. Therefore, we have $x_r(t) \in \mathcal{X}_r \subset \mathbb{R}^{24}$. Regarding the object states, we define o^{goal}
424 and o_t^{pen} as the goal and current pen orientations, which are both unit direction vectors. Then, we
425 transform o_t^{pen} to a new rotated coordinate frame that is constructed by o^{goal} being x axis $([1,0,0])$.
426 Note that the vector \bar{o}_t^{pen} after transformation is also a unit vector and it converges to x axis if
427 the pen is perfectly manipulated to goal orientation o^{goal} . In addition, we also include the cen-
428 ter of mass position p_t^{pen} and their corresponding velocities $\dot{p}_t^{\text{pen}}, \dot{o}_t^{\text{pen}}$ (all \mathbb{R}^3). Finally, we have
429 $x_o(t) = [p_t^{\text{pen}}, \bar{o}_t^{\text{pen}}, \dot{p}_t^{\text{pen}}, \dot{o}_t^{\text{pen}}] \in \mathcal{X}_o \subset \mathbb{R}^{12}$. As a result, $x(t)$ includes 36 states in total and we use
430 $T = 100$.

431 In this work, we only included the joint positions as the robot states (with the only exception of
432 NGF’s second-order policy) for the following reasons: 1) Given that these tasks are not repetitive,
433 we found that joint position information was sufficient to disambiguate the robot’s next action, 2)
434 even when ambiguity arises for a given joint position, object state information can help with disam-
435 biguation. Further, the impressive performance achieved by KODex in our experiments support this
436 design choice. Indeed, KODex is agnostic to this specific state design. One can incorporate velocity
437 information into the robot state space without the need of any changes to the training procedure.

438 C Task Success Criteria

439 The task success criteria are listed below. The settings were the same as proposed in [8].

440 **Tool Use:** The task is considered successful if at last time step T , the Euclidean distance between
441 the final nail position and the goal nail position is smaller than 0.01.

442 **Door Opening:** The task is considered successful if at last time step T , the door opening angle is
443 larger than 1.35 rad.

444 **Object Relocation:** At each time step t , if $\sqrt{|p^{\text{target}} - p_t^{\text{ball}}|^2} < 0.10$, then we have $\rho(t) = 1$. The
445 task is considered successful if $\sum_{t=1}^T \rho(t) > 10$.

446 **In-hand Reorientation:** At each time step t , if $o^{\text{goal}} \cdot o_t^{\text{pen}} > 0.90$ ($o^{\text{goal}} \cdot o_t^{\text{pen}}$ measures orientation
447 similarity), then we have $\rho(t) = 1$. The task is considered successful if $\sum_{t=1}^T \rho(t) > 10$.

448 D Sampling Procedure

449 We describe the sampling procedure in this section. The sample distributions used for RL training
450 and demo collection were identical, as suggested in [8]. The out-of-distribution data were generated
451 to evaluate the zero-shot out-of-distribution generalizability of each policy.

452 **Tool Use:** We randomly sampled the nail heights (h) from a uniform distributions. Within distri-
453 bution: we used $h \in \mathcal{H} \sim \mathcal{U}(0.1, 0.25)$; Out of distribution: we used $h \in \mathcal{H} \sim \mathcal{U}(0.05, 0.1) \cup$
454 $\mathcal{U}(0.25, 0.3)$.

455 **Door Opening:** We randomly sampled the door positions (xyz) from uniform distributions. Within
456 distribution: we used $x \in \mathcal{X} \sim \mathcal{U}(-0.3, 0)$, $y \in \mathcal{Y} \sim \mathcal{U}(0.2, 0.35)$, and $z \in \mathcal{Z} \sim \mathcal{U}(0.252, 0.402)$;
457 Out of distribution: we used $y \in \mathcal{Y} \sim \mathcal{U}(0.15, 0.2) \cup \mathcal{U}(0.35, 0.4)$ (x, z remained unchanged).

458 **Object Relocation:** We randomly sampled the target positions (xyz) from uniform distributions.
459 Within distribution: we used $x \in \mathcal{X} \sim \mathcal{U}(-0.25, 0.25)$, $y \in \mathcal{Y} \sim \mathcal{U}(-0.25, 0.25)$, and
460 $z \in \mathcal{Z} \sim \mathcal{U}(0.15, 0.35)$; Out of distribution: we used $z \in \mathcal{Z} \sim \mathcal{U}(0.35, 0.40)$ (x, y remained
461 unchanged).

462 **In-hand Reorientation:** We randomly sampled the pitch (α) and yaw (β) angles of the goal
463 orientation from uniform distributions. Within distribution: we used $\alpha \in \mathcal{A} \sim \mathcal{U}(-1, 1)$ and
464 $\beta \in \mathcal{B} \sim \mathcal{U}(-1, 1)$; Out of distribution: we used $\{(\alpha, \beta) \in (\mathcal{A}, \mathcal{B}) \sim (\mathcal{U}(-1, 1.2)), \mathcal{U}(1, 1.2)) \cup$
465 $(\mathcal{U}(1, 1.2)), \mathcal{U}(-1.2, 1)) \cup (\mathcal{U}(-1.2, 1)), \mathcal{U}(-1.2, -1)) \cup (\mathcal{U}(-1.2, -1)), \mathcal{U}(-1, 1.2))\}$.

466 **E Policy Design**

467 We show the detailed policy design in this section. All the baseline policies were trained to minimize
 468 the trajectory reproduction error.

469 **KODex:** The representation of the system is given as: $\mathbf{x}_r = [x_r^1, x_r^2, \dots, x_r^n]$ and $\mathbf{x}_o =$
 470 $[x_o^1, x_o^2, \dots, x_o^m]$ and superscript is used to index states. The details of the state design for each
 471 task is provided in Appendices B. In experiments, the vector-valued lifting functions ψ_r and ψ_o in
 472 (8) were polynomial basis function defined as

$$\begin{aligned} \psi_r &= \{x_r^i x_r^j\} \cup \{(x_r^i)^3\} \text{ for } i, j = 1, \dots, n \\ \psi_o &= \{x_o^i x_o^j\} \cup \{(x_o^i)^2 (x_o^j)\} \text{ for } i, j = 1, \dots, m \end{aligned} \tag{11}$$

473 Note that $x_r^i x_r^j / x_r^j x_r^i$ only appears once in lifting functions (similar to $x_o^i x_o^j / x_o^j x_o^i$), and we ignore t
 474 as the lifting functions are the same across the time horizon.

475 The choice of lifting functions can be viewed as the hyper-parameter of KODex. We make this
 476 choice as inspired from [24] and experimental results also indicate its effectiveness. Through all
 477 the experiments, we stucked with the same set of lifting functions, which helped to relieve us from
 478 extensive efforts of tuning the hyper-parameters, e.g. network layer size, that were necessary for
 479 baseline policies as shown in Appendices F.

480 **Full-connected Neural Network (NN):** The first baseline is a feedforward network that ingests the
 481 states $\mathbf{x}(1)$ and iteratively produces the predictions $\mathbf{x}(t), t = 2, \dots, T$ via the rollout of a Multilayer
 482 Perceptron (MLP). The reference joint trajectories ($\mathbf{x}_r(t)$) are then used to execute the robot with
 483 the learned controller C . The significance of this baseline is to evaluate a policy that produces a
 484 high-dimensional motion without any additional structure.

485 **Long Short-Term Memory (LSTM):** We create an LSTM-based policy under the same input-
 486 output flow as the NN policy. We also apply two fully-connected layers between the task input-
 487 output and the input/hidden state of the LSTM network. Similarly, the same controller C is
 488 deployed to track the reference joint trajectory. LSTM networks are known to be beneficial to imita-
 489 tion learning [30] and suitable for sequential processing [37], e.g. motion generation. Therefore,
 490 we expect to evaluate the performance of the recurrent structures in these tasks.

491 **Neural Dynamic Policy (NDP):** The Neural Dynamic Policy [17] embeds desired dynamical struc-
 492 ture as a layer in neural networks. Specifically, the parameters of the second order Dynamics Motion
 493 Primitives (DMP) are predicted as outputs of the preceding layers (MLP in [17]). As a result, it al-
 494 lows the overall policy easily reason in the space of trajectories and can be utilized for learning from
 495 demonstration. We train an NDP policy following the imitation learning pipeline described in [17].
 496 For each task, given $\mathbf{x}(1)$, the neural network components in NDP generate the parameters of DMPs
 497 (radial basis functions (RBFs) in [17]), which are forward integrated to produce the reference joint
 498 trajectories for tracking.

499 **Neural Geometric Fabrics policy (NGF):** The Neural Geometric Fabrics [4], a structured pol-
 500 icy class, that enables efficient skill learning for dexterous manipulation from demonstrations by
 501 leveraging structures induced by Geometric Fabrics [38]. Geometric Fabrics is a stable class of
 502 the Riemannian Motion Policy (RMP) [39]. It has been demonstrated that NGF outperforms RMP
 503 in policy learning for dexterous manipulation task in [4]. The NGF policy is defined in the con-
 504 figuration space of the robot, which is composed of a geometric policy, a potential policy and a
 505 damping term. More specifically, the NGF policy is constructed as follows: (1) define a geometric
 506 policy pair $[\mathbf{M}, \pi]$ and a potential policy pair $[\mathbf{M}_f, \pi_f]$ in the configuration space \mathbf{q} , (2) energize
 507 the geometric policy (project orthogonal to the direction of motion with \mathbf{p}_e) to create a collection of
 508 energy-preserving paths (the Geometric Fabric), and (3) force the Geometric Fabric with a potential
 509 defined by $[\mathbf{M}_f, \pi_f]$ and damp via b applied along $\dot{\mathbf{q}}$, which ensures convergence to the potential’s
 510 minima. The potential policy π_f is the gradient of a function of position only. Note that we param-
 511 eterize the geometric policy pair $[\mathbf{M}, \pi]$, the potential policy pair $[\mathbf{M}_f, \pi_f]$, and the damping scalar
 512 b with MLP networks and learn them from demonstration data.

513 **F Optimizing baseline model size**

514 As described in Appendices E, we stuck with the same set of lifting functions for KODex and
 515 report the task success rate when we trained KODex on training set and tested it on validation set
 516 in Table. 1. However, for baselines, the hyper-parameters were selected through a set of ablation
 517 experiments for each task using the training set over three choices of model size, including small
 518 size, median size and large size. We generated five random seeds for parameter initialization per
 519 model size, per baseline, and per task, as all learning based baseline models are sensitive to param-
 520 eter initialization [23]. For each baseline policy, we report the mean and standard deviation of the
 521 task success rate on the validation set over five random seeds in Tables. 2-5.

522 Based on these results, we selectd the model size that offers the best performance in terms of task
 523 success rate. In addition, these results indicate that, unlike KODex, extensive hyper-parameter tun-
 524 ing and various trials on parameter initialization for baseline models are necessary. Note that we use
 l to denote $\dim(\mathbf{x}(t))$.

Table 1: Task success rate on validation set (KODex)

Tool	Door	Relocation	Reorientation
100.0%	96.0%	88.0%	62.0%

Table 2: Hyper-parameters on NN Network Sizes

Success Rate (%) \ Task	Tool	Door	Relocation	Reorientation
Model Size				
MLP: (32, 64, 32)	0.4(±0.8)	0.0(±0.0)	0.4(±0.8)	6.8(±3.9)
MLP: (64, 128, 64)	0.0(±0.0)	0.4(±0.8)	1.2(±2.4)	10.4(±6.6)
MLP: (128, 256, 128)	0.0(±0.0)	0.0(±0.0)	0.8(±1.6)	6.0(±1.5)

525

Table 3: Hyper-parameters on LSTM Network Sizes

Success Rate (%) \ Task	Tool	Door	Relocation	Reorientation
Model Size				
LSTM: 200 fc: (l , 100), (200, l)	28.8(±25.0)	87.6(±10.3)	7.6(±5.9)	56.4(±7.4)
LSTM: 250 fc: (l , 175), (250, l)	60.8(±36.6)	80.8(±24.5)	7.6(±7.5)	48.0(±17.0)
LSTM: 300 fc: (l , 250), (300, l)	44.8(±31.8)	82.0(±13.9)	16.4(±14.5)	54.0(±11.0)

Table 4: Hyper-parameters on NDP Network Sizes

Success Rate (%) \ Task	Tool	Door	Relocation	Reorientation
Model Size				
MLP: (32, 64, 32) 10 RBFs	0.0(±0.0)	8.0(±2.5)	30.0(±9.3)	57.2(±8.6)
MLP: (64, 128, 64) 20 RBFs	16.8(±29.8)	40.8(±8.1)	74.0(±4.9)	59.2(±6.5)
MLP: (128, 256, 128) 30 RBFs	18.4(±31.9)	66.0(±5.2)	79.2(±7.7)	62.4(±7.8)

Table 5: Hyper-parameters on NGF Network Sizes

Success Rate (%)	Task	Tool	Door	Relocation	Reorientation
Model Size					
MLP: (64, 32)		99.2(\pm 1.6)	87.2(\pm 12.0)	87.6(\pm 8.5)	77.6(\pm 2.3)
MLP: (128, 64)		100.0(\pm0.0)	90.0(\pm 5.9)	94.4(\pm 3.2)	72.4(\pm 4.5)
MLP: (256, 128)		83.6(\pm 20.1)	90.8(\pm4.3)	95.2(\pm1.6)	78.4(\pm3.4)

526 G Hyper-parameters for controller learning

527 The hyper-parameters we used to learn the inverse dynamics controller C for each task were the same as listed in Table. 6. Note that we use l_r to denote $\dim(x_r(t))$.

Table 6: Hyper-parameters on controller learning

Hidden Layer	Activation	Learning Rate	Iteration
($4l_r, 4l_r, 2l_r$)	ReLU	0.0001	300

528

529 H Zero-Shot Out-of-Distribution Generalization

530 We generated a new set of 10,000 out-of-distribution samples to evaluate how the policies that were
531 trained on 200 demonstrations generalize to unseen samples (see Appendices D for details on the
532 sampling procedure). In Fig. 5, we report the task success rates of each method trained on the 200
533 demonstrations and tested on the 10,000 out-of-distribution samples. In addition, we also report the
534 task success rate of the expert policy on the same 10,000 out-of-distribution samples to establish
535 a baseline. Perhaps unsurprisingly, none of the methods are able to consistently outperform the
536 expert policy in most tasks. We observe that KODex is able to outperform the four baselines in Tool
537 Use task. In the other tasks, the highly-structured NGF performs the best, and KODex’s performs
538 comparably to NDP and LSTM.

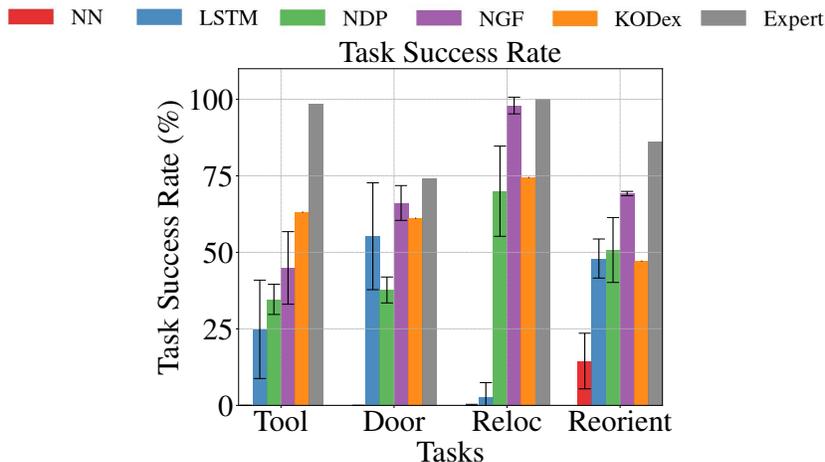


Figure 5: Zero-Shot Out-of-distribution task success rates

539 I Robustness to changes in physical properties

540 We evaluate the robustness of the reference dynamics learned by each method to changes in hand
541 mass or object mass for each task. This experiment is motivated by the fact that sim-to-real transfer
542 often involves changes in physical properties. Further, consistent use of robotic hardware could
543 result in changes to physical properties. Specifically, we consider four variations per task:

- 544 • Tool Use: i) *Heavy Object (Hammer)*: 0.25 (default) \rightarrow 0.85 (new), ii) *Light Object (Hammer)*:
545 0.25 (default) \rightarrow 0.10 (new), iii) *Light Hand (Palm)*: 4.0 (default) \rightarrow 1.0 (new), and iv) *Heavy*
546 *Hand (Palm)*: 4.0 (default) \rightarrow 8.0 (new)
- 547 • Door: i) *Heavy Object (Latch)*: 3.54 (default) \rightarrow 12.54 (new), ii) *Light Object (Latch)*: 3.54
548 (default) \rightarrow 0.54 (new), iii) *Light Hand (Palm)*: 4.0 (default) \rightarrow 1.5 (new), and iv) *Heavy Hand*
549 *(Palm)*: 4.0 (default) \rightarrow 7.0 (new)
- 550 • Relocation: i) *Heavy Object (Ball)*: 0.18 (default) \rightarrow 1.88 (new), ii) *Light Object (Ball)*: 0.18
551 (default) \rightarrow 0.05 (new), iii) *Light Hand (Palm)*: 4.0 (default) \rightarrow 3.0 (new), and iv) *Heavy Hand*
552 *(Palm)*: 4.0 (default) \rightarrow 5.0 (new);
- 553 • Reorientation: i) *Heavy Object (Pen)*: 1.5 (default) \rightarrow 9.5 (new), ii) *Light Object (Pen)*: 1.5
554 (default) \rightarrow 0.2 (new), iii) *Light Hand (Finger Knuckles)*: 0.008 (default) \rightarrow 0.0001 (new), and
555 iv) *Heavy Hand (Finger Knuckles)*: 0.008 (default) \rightarrow 0.20 (new)

556 It is important to note we held the reference dynamics learned by each method constant for this
557 experiment, irrespective of the changes to the hand or the object. Instead, we relearned the tracking
558 controller using 200 rollouts from the expert agent, following the procedure detailed in Section. 4.3.

559 In Tables. 7-10, we report the task success rate of KODex, and other baseline policies (all trained on
560 200 demonstrations) before and after relearning the controller. We also report the task success rates
561 of the expert agents to establish baselines.

562 We find that the Light Hand variation results in the lowest drop in performance across all methods
563 and all tasks, thus consequently relearning controllers does not offer any considerable improve-
564 ments. In contrast, all methods benefit from relearning the controller in the Heavy Hand variations,
565 as evidenced by the increased task success rates. Overall, we find that KODex outperforms all base-
566 lines, with the exception of NGF which performs better than KODex under a few variations and
567 tasks. Surprisingly, KODex (and some baselines) when used with the original controller outperform
568 the expert policy under a few variations (e.g., Heavy Object in Relocation task, and Heavy Object
569 in Door task). We believe this is due to the fact that KODex and the baselines learn to generate and
570 track desired trajectories separately, while the expert RL directly generates control inputs from state
571 information. In particular, the learned desired trajectories for a given tasks are likely invariant to
572 slight changes in physical properties. On rare occasions where this is not the case, we indeed find
573 that fine-tuning the tracking controllers worsens the performance.

574 These results demonstrate that changes to the robot/system dynamics can be handled by fine tuning
575 the tracking controller without the need for relearning the reference dynamics. Once again, KODex
576 is able to perform comparably to or outperform SOTA approaches despite its simplicity.

Table 7: Robustness to variations in the physical properties (Tool Use)

Success Rate Controller	Variation (%)	Heavy Object	Light Object	Light Hand	Heavy Hand
		Expert agent	93.5	66.2	65.4
KODex + Original controller		46.0	64.0	99.5	46.5
NN + Original controller		0.0(\pm 0.0)	0.0(\pm 0.0)	0.0(\pm 0.0)	0.7(\pm 1.4)
LSTM + Original controller		32.7(\pm 18.7)	35.0(\pm 22.1)	44.3(\pm 23.1)	52.7(\pm 27.5)
NDP + Original controller		0.0(\pm 0.0)	68.0(\pm 20.8)	45.4(\pm 37.4)	0.0(\pm 0.0)
NGF + Original controller		33.4(\pm 11.5)	62.9(\pm 27.5)	83.2(\pm 26.3)	40.3(\pm 20.2)
KODex + Expert-tuned controller		53.5	44.0	89.0	92.5
NN + Expert-tuned controller		0.0(\pm 0.0)	0.0(\pm 0.0)	0.2(\pm 0.4)	0.0(\pm 0.0)
LSTM + Expert-tuned controller		42.4(\pm 34.3)	33.7(\pm 14.9)	52.2(\pm 22.7)	69.9(\pm 19.4)
NDP + Expert-tuned controller		33.3(\pm 20.0)	23.8(\pm 24.4)	29.4(\pm 37.1)	39.8(\pm 24.5)
NGF + Expert-tuned controller		48.2(\pm 18.0)	48.7(\pm 12.2)	94.6(\pm 8.9)	82.1(\pm 7.5)

Table 8: Robustness to variations in the physical properties (Door)

Success Rate (%)	Variation	Heavy Object	Light Object	Light Hand	Heavy Hand
		Controller			
	Expert agent	45.2	91.7	82.0	74.9
	KODex + Original controller	57.0	97.0	56.5	33.5
	NN + Original controller	0.0(\pm 0.0)	0.2(\pm 0.4)	1.3(\pm 2.1)	0.0(\pm 0.0)
	LSTM + Original controller	34.4(\pm 8.7)	75.8(\pm 19.5)	38.1(\pm 10.4)	33.5(\pm 11.4)
	NDP + Original controller	22.1(\pm 1.9)	62.8(\pm 5.2)	51.1(\pm 4.9)	3.1(\pm 2.3)
	NGF + Original controller	48.7(\pm 6.7)	95.0(\pm 2.1)	42.1(\pm 11.0)	33.8(\pm 10.0)
	KODex + Expert-tuned controller	39.0	94.0	54.0	81.5
	NN + Expert-tuned controller	0.0(\pm 0.0)	0.0(\pm 0.0)	0.7(\pm 0.9)	0.0(\pm 0.0)
	LSTM + Expert-tuned controller	21.2(\pm 5.3)	75.4(\pm 18.0)	49.2(\pm 8.1)	56.9(\pm 18.7)
	NDP + Expert-tuned controller	15.5(\pm 3.0)	36.2(\pm 10.6)	25.5(\pm 4.4)	8.8(\pm 3.0)
	NGF + Expert-tuned controller	36.6(\pm 5.1)	95.5(\pm 1.8)	57.7(\pm 4.7)	77.1(\pm 6.7)

Table 9: Robustness to variations in the physical properties (Relocation)

Success Rate (%)	Variation	Heavy Object	Light Object	Light Hand	Heavy Hand
		Controller			
	Expert agent	77.0	100.0	100.0	100.0
	KODex + Original controller	19.5	89.5	82.5	21.5
	NN + Original controller	0.1(\pm 0.2)	1.6(\pm 2.5)	1.5(\pm 2.1)	1.7(\pm 2.2)
	LSTM + Original controller	0.4(\pm 0.4)	15.4(\pm 10.7)	9.5(\pm 8.1)	7.7(\pm 9.4)
	NDP + Original controller	13.5(\pm 5.0)	85.6(\pm 8.1)	72.1(\pm 9.6)	31.6(\pm 10.0)
	NGF + Original controller	25.8(\pm 4.9)	96.4(\pm 1.4)	96.6(\pm 0.97)	19.3(\pm 3.8)
	KODex + Expert-tuned controller	34.0	93.0	85.0	89.0
	NN + Expert-tuned controller	0.2(\pm 0.4)	0.6(\pm 0.7)	1.4(\pm 1.8)	1.5(\pm 2.3)
	LSTM + Expert-tuned controller	5.8(\pm 4.7)	15.2(\pm 12.5)	15.5(\pm 10.7)	14.1(\pm 9.3)
	NDP + Expert-tuned controller	19.9(\pm 5.8)	84.5(\pm 8.9)	63.2(\pm 15.0)	92.4(\pm 1.2)
	NGF + Expert-tuned controller	52.6(\pm 3.6)	98.1(\pm 1.2)	95.6(\pm 2.2)	94.5(\pm 0.9)

Table 10: Robustness to variations in the physical properties (Reorientation)

Success Rate (%)	Variation	Heavy Object	Light Object	Light Hand	Heavy Hand
		Controller			
	Expert agent	46.8	69.0	95.2	89.7
	KODex + Original controller	53.5	55.0	66.5	61.5
	NN + Original controller	4.7(\pm 2.6)	9.6(\pm 8.1)	9.5(\pm 6.4)	7.9(\pm 6.5)
	LSTM + Original controller	34.5(\pm 7.8)	52.3(\pm 10.6)	60.3(\pm 6.0)	55.6(\pm 7.8)
	NDP + Original controller	49.4(\pm 3.6)	58.4(\pm 6.4)	59.8(\pm 7.6)	55.7(\pm 9.7)
	NGF + Original controller	39.9(\pm 1.9)	57.1(\pm 2.2)	81.6(\pm 1.8)	73.4(\pm 3.8)
	KODex + Expert-tuned controller	52.0	63.0	71.5	65.5
	NN + Expert-tuned controller	1.5(\pm 0.9)	5.2(\pm 4.2)	3.8(\pm 1.7)	3.7(\pm 2.6)
	LSTM + Expert-tuned controller	43.5(\pm 7.9)	47.7(\pm 8.8)	61.4(\pm 4.2)	54.4(\pm 5.5)
	NDP + Expert-tuned controller	55.5(\pm 5.9)	59.0(\pm 5.5)	63.0(\pm 6.5)	57.0(\pm 7.5)
	NGF + Expert-tuned controller	49.1(\pm 2.6)	59.7(\pm 3.2)	79.4(\pm 1.9)	72.6(\pm 1.2)

577 J The impact of the choice of basis functions

578 We evaluate if KODex’s performance is impacted by different sets of polynomial functions that are
579 used as the lifting function. We trained all policies on 200 demos and tested them on 10,000 unseen
580 initial conditions.

581 **Design:** Specifically, we define four sets of observables (one of which was used in the original
582 submission). Let robot state: $x_r = [x_r^1, x_r^2, \dots, x_r^n]$ and $x_o = [x_o^1, x_o^2, \dots, x_o^m]$ denote the robot
583 and the object state, respectively, with superscript indexing the states. We then define four vector-
584 valued lifting functions ψ_r and ψ_o in (8) as follows

- 585 • Set 1

$$\psi_r = \{(x_r^i)^2\} \text{ for } i = 1, \dots, n$$

$$\psi_o = \{(x_o^i)^2\} \text{ for } i = 1, \dots, m$$

- 586 • Set 2

$$\psi_r = \{x_r^i x_r^j\} \text{ for } i, j = 1, \dots, n$$

$$\psi_o = \{x_o^i x_o^j\} \text{ for } i, j = 1, \dots, m$$

- 587 • **Set 3 (used in this work)**

$$\psi_r = \{x_r^i x_r^j\} \cup \{(x_r^i)^3\} \text{ for } i, j = 1, \dots, n$$

$$\psi_o = \{x_o^i x_o^j\} \cup \{(x_o^i)^2 (x_o^j)\} \text{ for } i, j = 1, \dots, m$$

- 588 • Set 4

$$\psi_r = \{x_r^i x_r^j\} \cup \{(x_r^i)^2 (x_r^j)\} \text{ for } i, j = 1, \dots, n$$

$$\psi_o = \{x_o^i x_o^j\} \cup \{(x_o^i)^2 (x_o^j)\} \text{ for } i, j = 1, \dots, m$$

We report the number of observables for each set and task combination in Table 11.

Table 11: Number of observables

Set \ Task	Tool	Door	Relocation	Reorientation
	n=26,m=15	n=28,m=7	n=30,m=12	n=24,m=12
Set 1	82	70	84	72
Set 2	512	469	585	414
Set 3 (ours)	763	546	759	582
Set 4	1413	1302	1629	1134

589

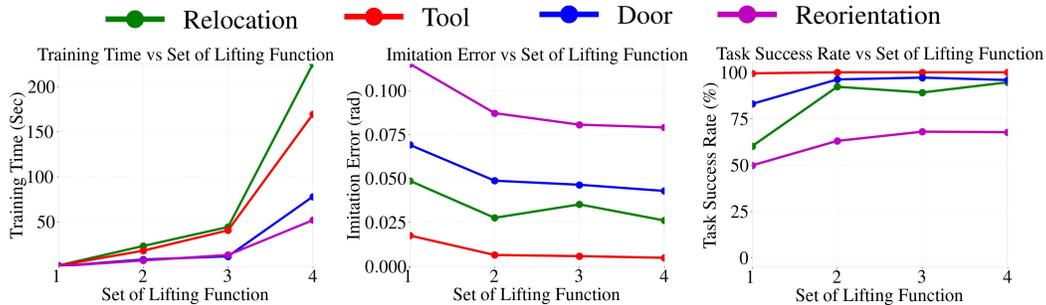


Figure 6: The effects of lifting function on training time (left), imitation error (center), and success rate (right).

590 **Discussion:** As shown in Fig. 6, it is clear that training time increases with the number of observ-
591 ables since the Moore–Penrose inverse requires more computation for higher-dimension matrices.
592 Importantly, KODex’s success rate across all tasks remained roughly the same for Sets 2, 3, and 4.
593 In general, as one would expect, increasing the number of observables tends to decrease imitation
594 error and increase task success rate. The only exception to this trend is observed for the Object
595 Relocation task, in which KODex performs marginally better when trained on Set 2 (585 observ-
596 ables) compared with it trained on Set 3 (759 observables). Taken together, these results suggest
597 that KODex’s performance is not highly sensitive to the specific choice of lifting function, as long
598 as sufficient expressivity is ensured.

599 **K Stability Analysis**

600 Another unique advantage of utilizing Koopman Operators to model the underlying dynamical sys-
 601 tem for dexterous manipulation tasks is that the learned policy is a linear dynamical system which
 602 can be readily inspected and analyzed, in stark contrast to SOTA methods built upon deep neural
 603 networks.

604 We analyzed the stability of the learned policy. For a linear dynamical system with complex conju-
 605 gate eigenvalues $\lambda_i = \theta_i \pm j\omega_i$, i.e., KODex with Koopman matrix \mathbf{K} , the system is asymptotically
 606 stable if all of the eigenvalues have magnitude ($\rho_i = \sqrt{\theta_i^2 + \omega_i^2}$) less than one. From the standpoint
 607 of control theory, it is beneficial to have an asymptotically stable system because of the guarantee
 608 that all system states will converge. However, from the standpoint of dexterous manipulation tasks
 609 considered in this work, strict stability might not be preferable because the final desired hand poses
 610 and object poses are not identical for different initial conditions. This represents a natural trade-off
 611 between safety and expressivity. As such, understanding how KODex addresses this trade-off can
 612 be illuminating.

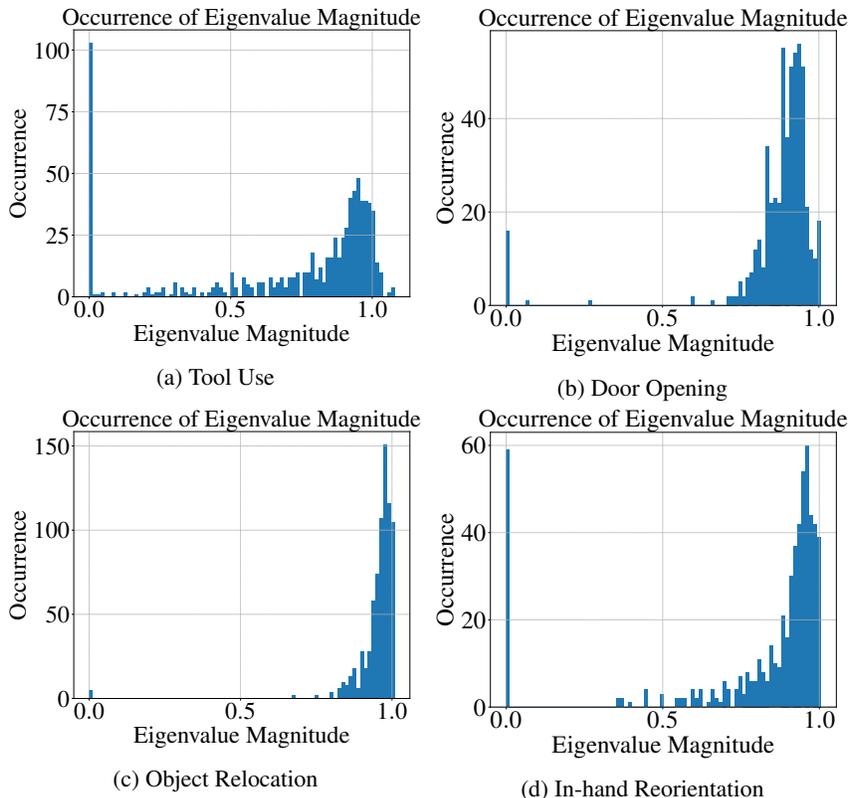


Figure 7: Occurrence of Eigenvalue Magnitude

Table 12: Maximun Eigenvalue Magnitude

Tool Use	Door Opening	Object Relocation	In-hand Reorientation
1.07888	1.00553	1.00859	1.00413

613 In Fig. 7, we report a histogram of the Koopman matrix’s eigenvalue magnitudes in each task. In
 614 addition, we report the maximum eigenvalue magnitude in Table. 12. Based on these results, we can
 615 see that i) most eigenvalues’ magnitudes are less than one, suggesting that KODex tends to learn
 616 nearly-stable policies that generate *safe trajectories* during execution, and ii) a few eigenvalues
 617 have magnitude larger than one, suggesting KODex does not prioritize stability, at the expense of
 618 expressivity required to achieve the reported performance.