

From Pixels to Torques with Linear Feedback

Jeong Hun Lee¹, Sam Schoedel¹, Aditya Bhardwaj², and Zachary Manchester¹

¹ Robotics Institute, Carnegie Mellon University, Pittsburgh 15213, USA
{jeonghunlee, sschoedel, zacm}@cmu.edu

² University of Chicago, Chicago 60637, USA
a7b@uchicago.edu

Abstract. We demonstrate the effectiveness of simple observer-based linear feedback policies for “pixels-to-torques” control of robotic systems using only a robot-facing camera. Specifically, we show that the matrices of an image-based Luenberger observer (linear state estimator) for a “student” output-feedback policy can be learned from demonstration data provided by a “teacher” state-feedback policy via simple linear-least-squares regression. The resulting linear output-feedback controller maps directly from high-dimensional raw images to torques while being amenable to the rich set of analytical tools from linear systems theory, allowing us to enforce closed-loop stability constraints in the learning problem. We also investigate a nonlinear extension of the method via the Koopman embedding. Finally, we demonstrate the surprising effectiveness of linear pixels-to-torques policies on a cartpole system, both in simulation and on real hardware. The policy successfully executes both stabilizing and swing-up trajectory-tracking tasks using only camera feedback while subject to model mismatch, process and sensor noise, perturbations, and occlusions. Open-source code for all experiments can be found here: https://roboticexplorationlab.org/projects/linear_pixels_to_torques.html

Keywords: Vision and Sensor-based Control · Data-Driven Control · Control Theory and Optimization

1 Introduction

Both model-based [25, 27, 32, 34, 43, 45, 47, 48, 51] and learning-based [1, 29, 36–38, 40, 46] control policies have demonstrated impressive results in robotics using vision-based sensory feedback. However, both suffer from inherent drawbacks: neural-network policies trained from images can require large amounts of data, which often must be generated in simulation before being transferred to a real robot. Meanwhile, model-based policies require separate machine-vision algorithms to extract features for downstream state estimation or model learning. Recently, data-driven switched linear control policies have been applied directly to image feedback with promising results [41]. However, this was limited to a quasi-static setting. We aim to apply data-driven linear feedback policies to the control of dynamic robotic systems directly from images as shown in Fig. 1.




Fig. 1: An overview of the problem we address, in which a simple linear output-feedback policy must control a robotic system using only feedback from images. In this work, the robotic system is both a simulated and real-world cartpole.

We demonstrate that it is possible to perform “pixels-to-torques” control of robotic systems with simple observer-based linear feedback policies in both simulation and on real hardware. Specifically, we show that an image-based linear output-feedback policy can be designed by learning a Luenberger observer (linear state estimator) via linear-least-squares (LLS) regression over supervised demonstration data. The result is a pixels-to-torques policy that is interpretable and can be analyzed using the rich set of tools from classical linear systems theory, all while inferring states directly from high-dimensional images. We leverage this interpretability to promote stability of the closed-loop policy via a convex extension of the LLS problem. The resulting linear output-feedback policy can achieve good performance while learning from small amounts of data, which we demonstrate on a classic cartpole system. Specifically, the policy is able to avoid sim-to-real transfer issues by learning directly from hardware data, and it successfully stabilizes the system while being robust to perturbations, unmodeled dynamics, and visual occlusions. To summarize, our contributions are:

- A data-efficient “student-teacher” methodology for learning linear, visual output-feedback policies.
- A convex formulation of the least-squares observer-learning problem that includes closed-loop stability constraints.
- A nonlinear extension of our method based on the Koopman embedding.
- Demonstration of linear “pixels-to-torques” control on a real-world cartpole.

The remainder of the paper is organized as follows: In Section 2, we provide a broad survey of related works on image-based control and estimation in robotics. In Section 3, we provide an overview of topics from linear control, state estimation, and Koopman-operator theory. In Section 4, we describe our linear policy framework as well as the student-teacher method developed to learn a Luenberger observer. We also introduce a Koopman-based extension for nonlinear control settings. Section 5 presents experimental results demonstrating the method’s performance on a cartpole system both in simulation and on real-world hardware. Lastly, we summarize our conclusions in Section 6.

2 Related Works

2.1 Learning Dynamics Models from Images

Over the last decade, the deep-learning community developed dynamics-model learning from images for model-based reinforcement learning (RL) and model-predictive control (MPC) [11, 12, 15, 20, 44, 49]. To avoid the high-dimensionality of pixel space, the dynamics were initially learned over lower-dimensional, object-centric states to explicitly track objects in the image scene [15, 49]. While the states in these works were explicitly predefined or incorporated into a known problem structure, subsequent work also learned the lower-dimensional latent state space for robotics settings, which include simulated gym environments [20, 44] and real-world object manipulation [12]. Specifically, Finn et. al. [12] and Watter et. al. [44] learned latent spaces where the learned dynamics were also locally linear and time varying. This allowed for simple stochastic policies to be effectively used for downstream control tasks. In contrast to object-centric models, Finn et. al. [11] also learned image-to-image dynamics (termed “visual foresight”) and coupled it with a stochastic MPC controller for pushing manipulation tasks.

Recently, “deep” Koopman methods have used deep neural networks to *lift* observations into a higher-dimensional latent space where the dynamics behave linearly [25, 33, 34, 47]. Laferriere et. al [25] and Xiao et. al [47] extend deep Koopman to pixel-to-control tasks for robotics, including cartpole stabilization. Specifically, images were first embedded to lower-dimensional latent spaces using autoencoders before being lifted via multi-layer perceptrons (MLPs). Lyu et. al. [34] expands upon this further to use contrastive encoders for cartpole-swingup tasks. We note the similarities of deep Koopman to *time-varying*-latent-space-dynamics learning by Finn et. al. [12] and Watter et. al. [44]. While both methods aim to learn dynamics models for downstream controllers, our method aims to learn a linear state estimator for which images are observations.

2.2 Learning Policies with Image Feedback

Learning control policies directly from visual inputs has been a widely used practice in the deep-RL community [1, 29, 36, 37]. While initially performed within virtual environments [36, 37], deep-neural-network image-feedback policies have quickly been extended to real-world robotics tasks, including manipulation via diffusion policies [8] and quadruped locomotion over various terrains with ego-centric vision [1]. These methods, while demonstrating impressive results, are limited by the need for large amounts of training data. To scale to learning on real-world hardware, Levine et. al. [29] proposed a guided policy search method for efficiently learning an end-to-end pixels-to-torques policy for manipulation tasks. Specifically, “guiding” linear controllers with privileged full-state feedback are optimized and used to provide supervised training data for learning the pixels-to-torques policy. Rather than learning a deep end-to-end policy, our method learns a simple Luenberger observer from the guided supervised data to achieve image-to-state estimation for a downstream state-feedback controller.

2.3 Vision-Based Pose and State Estimation

Deep learning has been widely applied to vision-based pose and state estimation across many robotics applications [26, 28, 39, 48, 50, 52]. Applications include both external feature detection for a downstream controller to track [39, 48] and estimation of the robot’s own pose or state [26, 28, 50, 52]. A commonality across these learning-based methods is the need for large amounts of data, necessitating data synthesis and sim-to-real transfer. As a result, Lu et. al. looked to differentiable rendering for both robot pose and state estimation from images [30, 31]. Specifically, Lu et. al. rendered predicted images of the robot and calculated the corresponding Jacobians for an extended Kalman filter, resulting in online robot-state estimation from a single camera [31]. However, this method requires a differentiable rendering model of the robot to generate predicted observations. Therefore, we aim to determine if it’s possible to directly learn an image-based Luenberger observer instead of relying on an explicit observation model.

2.4 Effectiveness of Linearization

Over many decades, linear controllers and state estimators have proven to be powerful tools within the controls community thanks to their simplicity, ease of analysis, and often surprisingly good performance on nonlinear systems [22, 23]. Recently, under the label of state-space models (SSMs), linear models have also become popular in the deep-learning community [17–19]. Specifically, Gu et. al. [17, 19] demonstrated that structured linear models in the latent spaces of deep networks can model long-range dependencies found in many applications like speech recognition, language processing, and genomics [18]. The most recent implementations have produced performance matching or exceeding that of widely-adopted transformers while being much more computationally efficient [17]. We note the strong similarity of SSMs to deep-Koopman methods [33], specifically in their learning of linear dynamics in a latent space.

In addition to deep-learning methods, non-neural linear alternatives have also been shown to be surprisingly effective in robotics [5, 21, 31] and vision-based applications [32, 41]. By applying approximate Koopman operators [24], both Haggerty et. al. [21] and Bruder et. al. [5] were able to model highly nonlinear, high-dimensional soft manipulators with no neural components. Within the context of quasi-static pile manipulation, Suh et. al. [41] showed that switched linear visual-foresight models can outperform deep counterparts in both prediction error and closed-loop control performance. However, applying these model-learning methods directly to image-based observations of open-loop unstable dynamical systems can be difficult. Koopman liftings of high-dimensional images are computationally expensive, and linear visual-foresight models have been restricted to quasi-static systems. Therefore, we aim to determine if it is possible to avoid the development of a Koopman dynamics model by directly learning an observer-based linear feedback policy.

3 Background

This section provides a brief review of linear control, state estimation, and Koopman-operator theory. We refer the reader to the substantial existing literature on these topics for further details [2, 3, 7, 22, 24, 35].

3.1 Linear Control

We aim to describe our controlled robot with a discrete-time linear time-invariant dynamics model,

$$x_{k+1} = Ax_k + Bu_k, \quad (1)$$

where $x_{k+1}, x_k \in \mathbb{R}^n$ are the robot states at time steps $k+1$ and k respectively; $u_k \in \mathbb{R}^m$ are the control inputs (i.e., actions) at time step k ; and $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$ are the dynamics coefficient matrices.

Using the dynamics described by (1), we can design a linear state-feedback controller to drive the robot state to the origin:

$$u_k = -Kx_k, \quad (2)$$

where $K \in \mathbb{R}^{m \times n}$ is the controller gain matrix. As a particular example, K could be determined by solving an LQR problem using the Riccati equation [22].

To evaluate the stability of the controller, we substitute (2) into (1) to obtain an expression for the closed-loop dynamics:

$$x_{k+1} = (A - BK)x_k, \quad (3)$$

where we can see that x_k will converge to the origin if $A - BK$ is stable:

$$|\text{eigvals}[A - BK]| < 1, \quad (4)$$

where $|\text{eigvals}[A - BK]|$ are the magnitudes of the eigenvalues of the closed-loop dynamics.

3.2 Linear State Estimation

As seen in (2), the implementation of a linear state-feedback controller requires knowledge of the full state, x_k . However, this is rarely the case; we usually only have access to information from the robot's sensors. For now, we will assume that the sensors only provide partial state information described by the following observation model:

$$y_k = Cx_k, \quad (5)$$

where $y_k \in \mathbb{R}^l$ are the observations (e.g., sensor readings, images, etc.) at time step k and $C \in \mathbb{R}^{l \times n}$ is the observation matrix.

An observer can then be designed to first predict states using the dynamics, (1), before correcting the predictions using sensor observations to provide the final full-state estimates. This is the core idea behind a Luenberger observer (linear state estimator):

$$\hat{x}_{k+1} = \underbrace{A\hat{x}_k + Bu_k}_{\text{prediction}} + \underbrace{L(y_{k+1} - \hat{y}_{k+1})}_{\text{correction}}, \quad (6)$$

where \hat{x}_{k+1} is the *estimated* state output by the observer at time step $k+1$; \hat{y}_{k+1} are the corresponding *predicted* observations; y_{k+1} are the corresponding *real-world* observations (i.e., sensor measurements); and $L \in \mathbb{R}^{n \times l}$ is the observer gain. L is often designed with measurement noise and model uncertainty in mind, while the estimator's internal state \hat{x}_k maintains a memory of aggregated information from previous time steps [22]. As an example, a Kalman filter can be used to design L to optimally estimate \hat{x}_k when all uncertainties in the system correspond to additive Gaussian noise.

The state estimates, \hat{x}_k , from the Luenberger observer at time step k can be subsequently passed to a linear state-feedback controller to calculate the corresponding control inputs,

$$u_k = -K\hat{x}_k. \quad (7)$$

We refer to this combined controller-observer system as an *output-feedback policy* with an *internal* state, \hat{x} .

Similarly to K , L can also be calculated to drive the state estimation errors, $x_k - \hat{x}_k$, asymptotically to zero for a given linear feedback controller. We first express \hat{y}_{k+1} in terms of \hat{x}_k and u_k by substituting the prediction term of (6) into (5):

$$\hat{y}_{k+1} = C(A\hat{x}_k + Bu_k). \quad (8)$$

Then, we substitute (7) and (8) into (6) to further simplify the observer,

$$\hat{x}_{k+1} = (I - LC)(A - BK)\hat{x}_k + Ly_{k+1}. \quad (9)$$

This allows us to analyze the stability of the Luenberger observer in a similar manner to the linear state-feedback controller as described by (4). Specifically, we can express the condition for asymptotically converging state-estimation errors as:

$$|\text{eigvals}[(I - LC)(A - BK)]| < 1. \quad (10)$$

For clarity, we will define the matrix A^{LK} :

$$A^{LK} = (I - LC)(A - BK). \quad (11)$$

3.3 Koopman-Operator Theory

Koopman-operator theory has seen increased adoption in robotics in recent years [5, 6, 13, 21, 25, 47]. This is due to its ability to apply model-based linear (or bilinear) control directly to nonlinear systems, which we denote as

$$x_{k+1} = f(x_k, u_k). \quad (12)$$

Assuming these dynamics are control affine, the nonlinear system can be represented *exactly* by an infinite-dimensional bilinear system [42]. This bilinear Koopman model takes the form,

$$z_{k+1} = Az_k + Bu_k + \sum_{i=1}^m u_k^i C^i z_k, \quad (13)$$

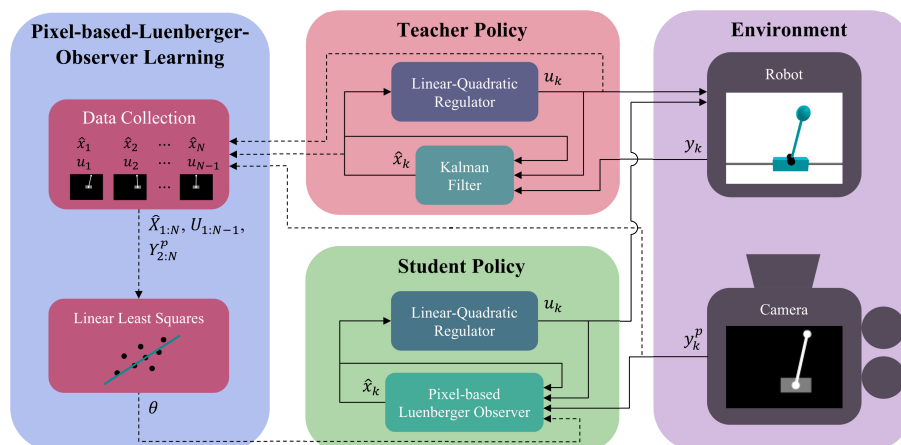


Fig. 2: Overview of learning a linear output-feedback policy for pixels-to-torques control. In this work, a linear-quadratic-Gaussian (LQG) “teacher” policy is designed to control a cartpole system. Teacher demonstration data is collected as trajectories of the robot’s estimated states, $\hat{X}_{1:N}$; control inputs, $U_{1:N-1}$; and corresponding images from a robot-facing camera, $Y_{2:N}^p$. Subsequent linear-least-squares (LLS) regression is performed over the data to determine the parameters, θ , of a “student” policy’s pixel-based Luenberger observer (linear state estimator). The student policy’s controller can also be learned separately or cloned from the teacher’s, as is the case in this work. The solid and dotted lines indicate processes that were performed online and offline respectively.

where z is the infinite-dimensional *lifting* of the robot states, x . In practice, this is approximated by a finite-dimensional nonlinear mapping over candidate basis functions [6, 13, 21] or deep neural networks [14, 25, 47]. We represent this finite-dimensional embedding as,

$$z = \phi(x), \quad (14)$$

where $z \in \mathbb{R}^j$. Likewise, ϕ is constructed in such a way that the “unlifting” is linear:

$$x = Gz. \quad (15)$$

4 Methodology

For a robot controlled via image-based feedback, the development of an appropriate observation model, (5), requires a differentiable renderer [30] that we may not have access to. In addition, performing a gain-calculation procedure, such as solving the Riccati equation, may be too computationally expensive to perform using dynamics and observation models in pixel space [41]. Therefore, we directly learn the parameters of a linear feedback policy’s Luenberger observer, which uses images from a robot-facing camera as sensory feedback. An overview of our method is shown in Fig. 2.

4.1 Pixel-Based Luenberger Observer

We begin by restating the Luenberger observer. Specifically, we substitute (8) into (6) to express the Luenberger observer in terms of just the predicted robot states, \hat{x} ; robot control inputs, u ; and output observations in the form of high-dimensional, raw pixel values from a robot-facing camera, y^p :

$$\hat{x}_{k+1} = \underbrace{(I - LC)A}_{A^L} \hat{x}_k + \underbrace{(I - LC)B}_{B^L} u_k + Ly_{k+1}^p, \quad (16)$$

where A^L , B^L denote new coefficient matrices. We additionally introduce an affine term, $d \in \mathbb{R}^n$, to handle linearizations about non-zero goal points:

$$\hat{x}_{k+1} = A^L \hat{x}_k + B^L u_k + Ly_{k+1}^p + d, \quad (17)$$

This yields the form of a pixel-based Luenberger observer whose coefficient matrices we will learn from subsequent supervised data. For simplification, we concatenate the matrices into

$$\theta = [A^L \ B^L \ L \ d]. \quad (18)$$

4.2 Learning a Pixel-Based Linear Output-Feedback Policy

To gather supervised data, we first leverage a predesigned “teacher” policy as shown in Fig. 2. In our case, the teacher is a privileged linear-quadratic regulator (LQR) policy that utilizes the robot’s built-in encoders. The teacher is used to collect supervised data in the form of demonstration trajectories of the predicted robot state, $\hat{X}_{1:N}$; control inputs, $U_{1:N-1}$; and pixel values from robot-facing-camera images, $Y_{2:N}^p$, which are used as observations by a subsequently learned linear output-feedback policy for pixels-to-torques control. We refer to this learned policy as the “student.”

The collected supervised trajectories are then concatenated and stored in the form of data matrices,

$$W = \begin{bmatrix} \hat{X}_{1:N-1} \\ U_{1:N-1} \\ Y_{2:N}^p \\ 1 \end{bmatrix} = \begin{bmatrix} \hat{x}_1 & \hat{x}_2 & \dots & \hat{x}_{N-1} \\ u_1 & u_2 & \dots & u_{N-1} \\ y_2^p & y_3^p & \dots & y_N^p \\ 1 & 1 & \dots & 1 \end{bmatrix}, \quad \hat{X}_{2:N} = [\hat{x}_2 \ \hat{x}_3 \ \dots \ \hat{x}_N]. \quad (19)$$

This allows us to formulate the learning of the student’s pixel-based Luenberger observer described by (17) as the following linear-least-squares (LLS) problem:

$$\underset{\theta}{\text{minimize}} \left\| \theta W - \hat{X}_{2:N} \right\|_2^2. \quad (20)$$

We additionally introduce sparsity-promoting L_1 regularization to reduce overfitting:

$$\underset{\theta}{\text{minimize}} \left\| \theta W - \hat{X}_{2:N} \right\|_2^2 + \lambda \|L\|_1. \quad (21)$$

By combining the teacher’s state-feedback controller with the newly designed Luenberger observer, the student output-feedback policy can perform closed-loop pixels-to-torques control as shown in Fig. 2. We note that regression can be

separately performed over states and control histories to also learn a gain, K , for a new student controller. However, if the teacher is a linear feedback policy, this is equivalent to using the teacher’s controller.

4.3 Enforcing Stability with a Linear-Matrix-Inequality

By leveraging existing techniques from linear systems theory [4, 22], we introduce a convex constraint to the observer-learning problem to ensure that state-estimation errors are guaranteed not to diverge. Specifically, we note that the stability condition described by (10) is equivalent to the following spectral-norm condition,

$$\|A^{LK}\|_2 < 1, \quad (22)$$

which can be rewritten as the following linear-matrix-inequality (LMI) [4]:

$$\begin{bmatrix} I & A^{LK} \\ A^{LK} & I \end{bmatrix} \succ 0, \quad (23)$$

where \succ indicates positive definiteness. Enforcing this constraint in the LLS problem requires privileged information of the teacher’s linear controller gain, K , which we assume access to. Combining (21) with (23) extends the original LLS problem into the following convex program, which we solve using the semidefinite programming solver COSMO [16]:

$$\begin{aligned} & \underset{\theta}{\text{minimize}} && \left\| \theta W - \hat{X}_{2:N} \right\|_2^2 + \lambda \|L\|_1, \\ & \text{subject to} && \begin{bmatrix} I & A^{LK} \\ A^{LK} & I \end{bmatrix} \succ 0, \end{aligned} \quad (24)$$

where θ and W are now defined as

$$W = \begin{bmatrix} \hat{X}_{1:N-1} \\ Y_{2:N}^p \\ 1 \end{bmatrix}, \quad \theta = [A^{LK} \ L \ d].$$

4.4 Extension to Nonlinear Systems via Koopman Embedding

For nonlinear control and estimation, the nonlinear dynamics can be linearized about a reference trajectory before being tracked by a time-varying linear controller (e.g., time-varying LQR) and estimator. However, doing so requires learning gains for every time step, K_k and L_k , which can become computationally expensive and data-inefficient with high-dimensional image-based observations. Therefore, we introduce a Koopman-based extension to the pixel-based linear output-feedback policy that only requires learning a single Luenberger gain, L .

To do so, we start with a pre-specified choice of Koopman embedding and corresponding unlifting,

$$z_k = \phi(x_k), \quad (25)$$

$$x_k = Gz_k. \quad (26)$$

We then replace the linear dynamics (1) with the bilinear Koopman dynamics (13) in our Luenberger Observer (6). Doing so yields the following Luenberger observer w.r.t. z :

$$z_{k+1} = Az_k + Bu_k + \sum_{i=1}^m u_k^i C^i z_k + L(y_k - \hat{y}_k), \quad (27)$$

which can be further simplified in a similar manner to (16). An additional affine term can also be introduced, yielding:

$$z_{k+1} = A^L z_k + B^L u_k + \sum_{i=1}^m u_k^i C^{L,i} z_k + Ly_{k+1} + d. \quad (28)$$

We may now formulate a new LLS problem in the same manner as (21):

$$\underset{\theta}{\text{minimize}} \left\| \theta W - \hat{Z}_{2:N} \right\|_2^2 + \lambda \|\theta\|_1, \quad (29)$$

where $\hat{Z}_{2:N}$ and W are now defined as

$$\hat{Z}_{2:N} = [\phi(\hat{x}_2) \ \phi(\hat{x}_3) \ \dots \ \phi(\hat{x}_N)],$$

$$W = \begin{bmatrix} \hat{Z}_{1:N-1} \\ U_{1:N-1} \\ (U^1 \hat{Z})_{1:N-1} \\ \vdots \\ (U^m \hat{Z})_{1:N-1} \\ Y_{2:N}^p \\ 1 \end{bmatrix} = \begin{bmatrix} \hat{z}_1 & \hat{z}_2 & \dots & \hat{z}_{N-1} \\ u_1 & u_2 & \dots & u_{N-1} \\ u_1^1 \hat{z}_1 & u_2^1 \hat{z}_2 & \dots & u_{N-1}^1 \hat{z}_{N-1} \\ \vdots & \vdots & \ddots & \vdots \\ u_1^m \hat{z}_1 & u_2^m \hat{z}_2 & \dots & u_{N-1}^m \hat{z}_{N-1} \\ y_2^p & y_3^p & \dots & y_N^p \\ 1 & 1 & \dots & 1 \end{bmatrix},$$

and θ represents the concatenated matrices of the Koopman Luenberger observer:

$$\theta = [A^L \ B^L \ C^{L,1} \ \dots \ C^{L,m} \ L \ d].$$


As before, this Koopman-based extension of the pixel-based Luenberger observer can be combined with the teacher’s state-feedback controller (e.g., time-varying LQR) or a newly designed Koopman-based MPC controller [6, 13, 24]. The lifting and unlifting operations, (25)-(26), can be used to pass state information between a state-feedback controller and the Koopman-based Luenberger observer. This creates a pixels-to-torques policy that can now track trajectories for nonlinear systems.

5 Experimental Results


5.1 Simulation Setup

We evaluate the closed-loop performance of our pixel-based linear output-feedback policy in a simulated cartpole environment. We specify two models: a *nominal* model, which is simplified and contains up to 5% parametric model error, and a *true* model, which is used exclusively for simulating the system in a “real-world” environment to evaluate the algorithm’s performance. We additionally

introduce Gaussian noise applied to the control inputs of the true model. The nominal model is used to design the teacher LQG policy, which uses encoder measurements of the cartpole’s configuration as observations for its state estimator. Observation noise is simulated as encoder quantization in the true model. To simulate a camera sensor for the pixels-to-torques student policy, we use Makie [9] to render greyscale images of the simulated cartpole at a resolution of 125×160 pixels. To discern these image-based observations from the true, “real-world” system, we visualize the true system in Meshcat [10]. The pixels-to-torques student policy is learned on a computer equipped with a 64-core AMD Threadripper CPU and 64GB of RAM.



(a) A visualization of a cartpole-stabilization trajectory executed by a pixel-based linear output-feedback policy. The top row depicts the simulated cartpole environment while the bottom row shows the corresponding rendered images (i.e., simulated camera) used by the pixels-to-torques policy. The goal configuration is shown in red.



(b) Stabilization-error performance

(c) Success rate of stabilization

Fig. 3: Stabilization performance vs. training trajectories of a pixel-based linear output-feedback policy tasked with stabilizing a cartpole using a robot-facing camera. Test stabilizations from 100 different initial conditions were evaluated with stabilization error defined as the L^2 error of the final state w.r.t the upright goal state. The median error is shown as a thick line, while the shaded regions represent the 5% to 95% bounds. The corresponding success rate of stabilization from the 100 initial conditions is also shown.





Fig. 4: Heat-map visualizations of each normalized row of the pixel-based Luenberger observer’s gain matrix, L . A cartpole visualization is also overlaid for reference. Each row of L corresponds to the correction an image observation contributes to a respective state variable. Interestingly, visual features can be distinguished for each state variable: the cart for the cart position with the addition of the pole tip for the pole angle. The velocities also have similar features.

5.2 Simulation Results


We evaluate the data-efficiency of learning the pixel-based linear output-feedback policy on cartpole-stabilization tasks as shown in Fig. 3. Specifically, we evaluate the policy’s stabilization performance across 100 initial conditions chosen from the same distribution as the training demonstrations. Each training demonstration consists of 150 samples over 2.5 seconds. As shown in Fig. 3b and Fig. 3c, the learned pixels-to-torques policy converges to a small stabilizing error with a 97% success rate with only 40 training trajectories. In addition, the student policy achieves a 100% success rate when trained with only 80 trajectories. Success is defined as the cart position being within 17.6 cm (half the pole length) of the origin with a pole angle within 2° of upright.

As shown in Fig. 4, we also investigate the observer gain of the pixel-based linear feedback policy. Specifically, we visualize each row of the policy’s observer gain, L , as a normalized heat map to discern the pixels that contribute most to the respective state variable’s estimation. Interestingly, doing so yields identifiable visual features: Pixels that correspond to the cart contribute most to the cart-position estimate while the pole-tip pixels additionally contribute to the pole-angle estimate. In addition, the sign of the gain correctly corresponds to the physical location of the respective pixel.

We additionally showcase the Koopman-based extension of the pixel-based linear output-feedback policy for tracking a nonlinear system as shown in Fig. 5. Specifically, we demonstrate the Koopman-based policy’s ability to track a reference swing-up trajectory on the cartpole. The 76-dimensional Koopman lifting consists of the cartpole states, 4th-order Fourier features, and 6th-order Chebyshev polynomial features. The reference trajectory was designed using iterative LQR before being tracked by a time-varying-LQG teacher policy for data collection. The Koopman-based policy is trained on 20 demonstrations, each consisting of 75 samples over 2.5 seconds. As shown in Fig. 5b, the policy is able to properly track the swing-up trajectory despite model mismatch and process noise.



(a) Cartpole swing-up trajectory visualization. The goal configuration is shown in red.



(b) Time histories of the cartpole configurations and corresponding control inputs. The ground truth values and goal-reference trajectory are also shown.

Fig. 5: A successful cartpole swing up performed by a Koopman-based extension of the pixel-based linear output-feedback policy. The policy is able to overcome process noise and model mismatch to track a reference trajectory on a nonlinear system.

5.3 Hardware Setup

We develop a physical cartpole system, shown in Fig. 6, to evaluate the efficacy of the pixel-based linear output-feedback policy on real-world hardware. The cartpole itself consists of a single brushless motor, a cart that rides along an aluminum track, and a pole laser cut from white acrylic. A black background is added to aid image contrast, and reflective components are also covered with a black cloth during runtime. Opposing the cartpole is a Jetson Nano and Arducam OV2311 camera mounted on a tripod. The camera has a global shutter to

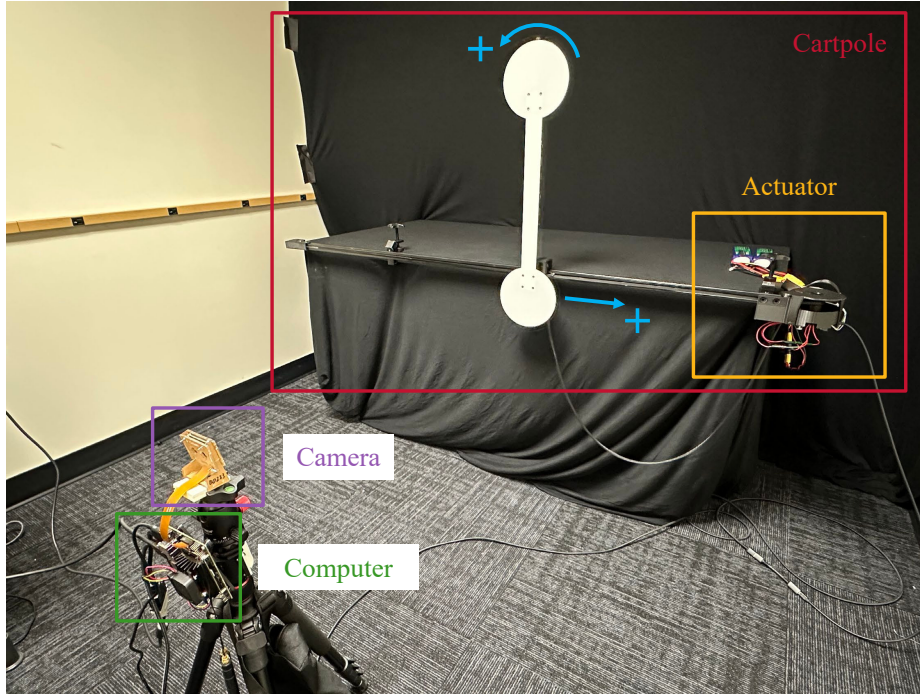




Fig. 6: Overview of the cartpole hardware used for real-world demonstrations of pixels-to-torques control performed by linear output-feedback policies. The cartpole hardware is actuated by a brushless motor while an opposing global-shutter Arducam OV2311 camera captures images of the cartpole. A Jetson Nano acts as the primary computer tasked with processing the images and executing the pixel-based student policy (or corresponding teacher) before sending the control inputs to the cartpole motor. This operation is performed at 60 Hz to match the camera’s maximum frame rate. Arrows indicate positive directions for the cart position, pole angle, and corresponding velocities.

avoid rolling-shutter distortion. The Jetson Nano acts as the main computer to perform online tasks, which consists of basic image processing (Gaussian blurring and Otsu thresholding) before executing the pixel-based linear output-feedback policy for state estimation and control. The system then sends control inputs to the motor at 60 Hz to match the frame rate of the camera. Image capture and processing is performed with OpenCV at an image resolution of 115×140 pixels. The pixels-to-torques student policy is trained with only 20 stabilizing demonstrations on the hardware, each with 150 samples over 2.5 seconds. The learning process is performed on a separate computer equipped with a 64-core AMD Threadripper CPU and 64GB of RAM, resulting in a solve time of about 4 minutes and 25 seconds.



(a) Snapshots of a real-world cartpole being stabilized by a pixel-based linear output-feedback policy when perturbed by a force, F , at times $t = 0$ and $t = 5.5$ seconds. The top row shows unfiltered images of the cartpole system taken by a separate cellphone camera while the bottom row shows the corresponding Arducam-camera images that are sent to the policy's Luenberger observer.



(b) Time histories of the cartpole configurations and corresponding control inputs. The teacher's estimates and anticipated control inputs are also shown as a reference.

Fig. 7: A hardware demonstration of a pixel-based linear output-feedback policy stabilizing a real-world cartpole based on image-feedback from an Arducam OV2311 camera. Lightweight processing is performed on the camera images in the form of Gaussian blur and Otsu thresholding before being sent to the policy's pixel-based Luenberger observer. Surprisingly, the pixel-based linear output-feedback policy is able to successfully stabilize the cartpole even under the presence of large perturbations and visual occlusions.

5.4 Hardware Results

As shown in Fig. 7, the pixels-to-torques policy is able to successfully recover from perturbations and stabilize the real-world cartpole. Surprisingly, the policy is also robust to major occlusions, as shown in Fig. 7a. While the occlusions do cause decreased state-estimation performance, as shown in Fig. 7b, the policy’s state-feedback controller is still able to overcome the disturbance to stabilize the cartpole. We attribute this to the sparse representation of L as depicted in Fig. 4, which prevents a majority of the occluded pixels from affecting the policy’s internal state estimate.

6 Conclusions

We have shown that it is possible to perform “pixels-to-torques” control of a highly dynamic robotic system with simple linear output-feedback policies. These pixel-based linear policies are amenable to analysis via linear systems theory while offering surprisingly effective control in the presence of model mismatch, disturbances, and visual occlusions, as demonstrated on a real-world cartpole system. We additionally introduced a Koopman-based extension of the pixels-to-torques policy for nonlinear systems.

Our linear pixels-to-torques approach has several limitations: First, the teacher-student methodology assumes that a successful teacher policy can be designed a priori with an appropriate internal state representation. Second, the method assumes un-obstructed observations during data collection. Third, by mapping raw pixel values directly to state estimates, our approach may also be sensitive to the camera’s calibration and pose. Finally, the nonlinear extension of our method also suffers from the limitations of Koopman-based approaches more broadly, which include the difficulty of crafting a good set of features or “observables” for lifting.

Several interesting directions for future work remain: It should be possible to apply linear visual feedback to systems with egocentric and/or multiple cameras. There are also several ways of combining our method with adaptive control or online learning techniques. Finally, we also plan to further explore the rich connection between pixel-based output feedback, Koopman operators, and diffusion policies.

Acknowledgments. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE2140739.

References

1. Agarwal, A., Kumar, A., Malik, J., Pathak, D.: Legged locomotion in challenging terrains using egocentric vision. In: Conference on Robot Learning. pp. 403–415. PMLR (2023)
2. Bertsekas, D.: Reinforcement learning and optimal control. Athena Scientific (2019)

3. Bertsekas, D., Shreve, S.E.: Stochastic optimal control: the discrete-time case, vol. 5. Athena Scientific (1996)
4. Boyd, S., El Ghaoui, L., Feron, E., Balakrishnan, V.: Linear Matrix Inequalities in System and Control Theory. Society for Industrial and Applied Mathematics (1994). <https://doi.org/10.1137/1.9781611970777>, <https://epubs.siam.org/doi/abs/10.1137/1.9781611970777>, [_eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9781611970777](https://epubs.siam.org/doi/pdf/10.1137/1.9781611970777)
5. Bruder, D., Fu, X., Gillespie, R.B., Remy, C.D., Vasudevan, R.: Data-Driven Control of Soft Robots Using Koopman Operator Theory. *IEEE Transactions on Robotics* **37**(3), 948–961 (2021). <https://doi.org/10.1109/TR0.2020.3038693>
6. Bruder, D., Fu, X., Vasudevan, R.: Advantages of Bilinear Koopman Realizations for the Modeling and Control of Systems With Unknown Dynamics. *IEEE Robotics and Automation Letters* **6**(3), 4369–4376 (2021). <https://doi.org/10.1109/LRA.2021.3068117>
7. Brunton, S.L., Budišić, M., Kaiser, E., Kutz, J.N.: Modern Koopman Theory for Dynamical Systems (2021)
8. Chi, C., Xu, Z., Feng, S., Cousineau, E., Du, Y., Burchfiel, B., Tedrake, R., Song, S.: Diffusion Policy: Visuomotor Policy Learning via Action Diffusion (2024), [_eprint: 2303.04137](https://arxiv.org/abs/2303.04137)
9. Danisch, S., Krumbiegel, J.: Makie.jl: Flexible high-performance data visualization for Julia. *Journal of Open Source Software* **6**(65), 3349 (2021). <https://doi.org/10.21105/joss.03349>, <https://doi.org/10.21105/joss.03349>
10. Deits, R.: MeshCat.jl (Apr 2024), <https://github.com/rdeits/MeshCat.jl>
11. Finn, C., Levine, S.: Deep visual foresight for planning robot motion. In: 2017 IEEE International Conference on Robotics and Automation (ICRA). pp. 2786–2793 (2017). <https://doi.org/10.1109/ICRA.2017.7989324>
12. Finn, C., Tan, X.Y., Duan, Y., Darrell, T., Levine, S., Abbeel, P.: Deep spatial autoencoders for visuomotor learning. In: 2016 IEEE International Conference on Robotics and Automation (ICRA). pp. 512–519 (2016). <https://doi.org/10.1109/ICRA.2016.7487173>
13. Folkestad, C., Burdick, J.W.: Koopman NMPC: Koopman-based learning and nonlinear model predictive control of control-affine systems. In: 2021 IEEE International Conference on Robotics and Automation (ICRA). pp. 7350–7356. IEEE (2021)
14. Folkestad, C., Wei, S.X., Burdick, J.W.: KoopNet: Joint Learning of Koopman Bilinear Models and Function Dictionaries with Application to Quadrotor Trajectory Tracking. In: 2022 International Conference on Robotics and Automation (ICRA). pp. 1344–1350 (2022). <https://doi.org/10.1109/ICRA46639.2022.9811896>
15. Fragkiadaki, K., Agrawal, P., Levine, S., Malik, J.: Learning Visual Predictive Models of Physics for Playing Billiards (2016), [_eprint: 1511.07404](https://arxiv.org/abs/1511.07404)
16. Garstka, M., Cannon, M., Goulart, P.: COSMO: A Conic Operator Splitting Method for Convex Conic Problems. *Journal of Optimization Theory and Applications* **190**(3), 779–810 (2021). <https://doi.org/10.1007/s10957-021-01896-x>, <https://doi.org/10.1007/s10957-021-01896-x>, publisher: Springer
17. Gu, A., Dao, T.: Mamba: Linear-Time Sequence Modeling with Selective State Spaces (2023), [_eprint: 2312.00752](https://arxiv.org/abs/2312.00752)
18. Gu, A., Dao, T., Ermon, S., Rudra, A., Ré, C.: Hippo: Recurrent memory with optimal polynomial projections. *Advances in neural information processing systems* **33**, 1474–1487 (2020)
19. Gu, A., Goel, K., Ré, C.: Efficiently Modeling Long Sequences with Structured State Spaces (2022), [_eprint: 2111.00396](https://arxiv.org/abs/2111.00396)

20. Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., Davidson, J.: Learning Latent Dynamics for Planning from Pixels (2019), [_eprint: 1811.04551](#)
21. Haggerty, D.A., Banks, M.J., Kamenar, E., Cao, A.B., Curtis, P.C., Mezić, I., Hawkes, E.W.: Control of soft robots with inertial dynamics. *Science Robotics* **8**(81), eadd6864 (2023). <https://doi.org/10.1126/scirobotics.add6864>, <https://www.science.org/doi/abs/10.1126/scirobotics.add6864>, [_eprint: https://www.science.org/doi/pdf/10.1126/scirobotics.add6864](https://www.science.org/doi/pdf/10.1126/scirobotics.add6864)
22. Kailath, T.: *Linear systems*, vol. 156. Prentice-Hall Englewood Cliffs, NJ (1980)
23. Kalman, R.E.: A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering* **82**(1), 35–45 (Mar 1960). <https://doi.org/10.1115/1.3662552>, <https://doi.org/10.1115/1.3662552>
24. Korda, M., Mezić, I.: Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *Automatica* **93**, 149–160 (2018). <https://doi.org/10.1016/j.automatica.2018.03.046>, <https://www.sciencedirect.com/science/article/pii/S000510981830133X>
25. Laferrière, P., Laferrière, S., Dahdah, S., Forbes, J.R., Paull, L.: Deep Koopman Representation for Control over Images (DKRCI). In: 2021 18th Conference on Robots and Vision (CRV). pp. 158–164 (2021). <https://doi.org/10.1109/CRV52889.2021.00029>
26. Lambrecht, J., Kästner, L.: Towards the Usage of Synthetic Data for Marker-Less Pose Estimation of Articulated Robots in RGB Images. In: 2019 19th International Conference on Advanced Robotics (ICAR). pp. 240–247 (2019). <https://doi.org/10.1109/ICAR46387.2019.8981600>
27. Larsen, P., Patience, D., Rawlings, J.: Industrial crystallization process control. *IEEE Control Systems Magazine* **26**(4), 70–80 (2006). <https://doi.org/10.1109/MCS.2006.1657878>
28. Lee, T.E., Tremblay, J., To, T., Cheng, J., Mosier, T., Kroemer, O., Fox, D., Birchfield, S.: Camera-to-robot pose estimation from a single image. In: 2020 IEEE International Conference on Robotics and Automation (ICRA). pp. 9426–9432. IEEE (2020)
29. Levine, S., Finn, C., Darrell, T., Abbeel, P.: End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research* **17**(1), 1334–1373 (2016), publisher: JMLR. org
30. Lu, J., Liu, F., Girerd, C., Yip, M.C.: Image-based Pose Estimation and Shape Reconstruction for Robot Manipulators and Soft, Continuum Robots via Differentiable Rendering. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 560–567 (2023). <https://doi.org/10.1109/ICRA48891.2023.10161066>
31. Lu, J., Richter, F., Lin, S., Yip, M.C.: Tracking Snake-like Robots in the Wild Using Only a Single Camera (2023), [_eprint: 2309.15700](#)
32. Lu, Q., Zavala, V.M.: Image-Based Model Predictive Control via Dynamic Mode Decomposition (2021), [_eprint: 2006.06727](#)
33. Lusch, B., Kutz, J.N., Brunton, S.L.: Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications* **9**(1), 4950 (Nov 2018). <https://doi.org/10.1038/s41467-018-07210-0>, <https://doi.org/10.1038/s41467-018-07210-0>
34. Lyu, X., Hu, H., Siriya, S., Pu, Y., Chen, M.: Task-Oriented Koopman-Based Control with Contrastive Encoder. In: 7th Annual Conference on Robot Learning (2023), <https://openreview.net/forum?id=q0VAoefCI2>
35. Mezić, I.: Koopman operator, geometry, and learning of dynamical systems. *Not. Am. Math. Soc.* **68**(7), 1087–1105 (2021)

36. Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M.: Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602 (2013)
37. Pathak, D., Agrawal, P., Efros, A.A., Darrell, T.: Curiosity-driven exploration by self-supervised prediction. In: International conference on machine learning. pp. 2778–2787. PMLR (2017)
38. Quillen, D., Jang, E., Nachum, O., Finn, C., Ibarz, J., Levine, S.: Deep Reinforcement Learning for Vision-Based Robotic Grasping: A Simulated Comparative Evaluation of Off-Policy Methods. In: 2018 IEEE International Conference on Robotics and Automation (ICRA). pp. 6284–6291 (2018). <https://doi.org/10.1109/ICRA.2018.8461039>
39. Richter, F., Shen, S., Liu, F., Huang, J., Funk, E.K., Orosco, R.K., Yip, M.C.: Autonomous Robotic Suction to Clear the Surgical Field for Hemostasis Using Image-Based Blood Flow Detection. *IEEE Robotics and Automation Letters* **6**(2), 1383–1390 (2021). <https://doi.org/10.1109/LRA.2021.3056057>
40. Shi, L.X., Sharma, A., Zhao, T.Z., Finn, C.: Waypoint-Based Imitation Learning for Robotic Manipulation. In: Tan, J., Toussaint, M., Darvish, K. (eds.) Proceedings of The 7th Conference on Robot Learning. Proceedings of Machine Learning Research, vol. 229, pp. 2195–2209. PMLR (Nov 2023), <https://proceedings.mlr.press/v229/shi23b.html>
41. Suh, H.T., Tedrake, R.: The surprising effectiveness of linear models for visual foresight in object pile manipulation. In: Algorithmic Foundations of Robotics XIV: Proceedings of the Fourteenth Workshop on the Algorithmic Foundations of Robotics 14. pp. 347–363. Springer (2021)
42. Surana, A.: Koopman operator based observer synthesis for control-affine nonlinear systems. In: 2016 IEEE 55th Conference on Decision and Control (CDC). pp. 6492–6499 (2016). <https://doi.org/10.1109/CDC.2016.7799268>
43. Villarreal, O., Barasuol, V., Wensing, P.M., Caldwell, D.G., Semini, C.: MPC-based controller with terrain insight for dynamic legged locomotion. In: 2020 IEEE International Conference on Robotics and Automation (ICRA). pp. 2436–2442. IEEE (2020)
44. Watter, M., Springenberg, J.T., Boedecker, J., Riedmiller, M.: Embed to Control: A Locally Linear Latent Dynamics Model for Control from Raw Images (2015), [_eprint: 1506.07365](https://arxiv.org/abs/1506.07365)
45. Wilson, W., Williams Hulls, C., Bell, G.: Relative end-effector control using Cartesian position based visual servoing. *IEEE Transactions on Robotics and Automation* **12**(5), 684–696 (1996). <https://doi.org/10.1109/70.538974>
46. Wu, B., Nair, S., Fei-Fei, L., Finn, C.: Example-Driven Model-Based Reinforcement Learning for Solving Long-Horizon Visuomotor Tasks (2022), [_eprint: 2109.10312](https://arxiv.org/abs/2109.10312)
47. Xiao, Y., Tang, Z., Xu, X., Zhang, X., Shi, Y.: A deep Koopman operator-based modelling approach for long-term prediction of dynamics with pixel-level measurements. *CAAI Transactions on Intelligence Technology* **9**(1), 178–196 (Feb 2023). <https://doi.org/10.1049/cit2.12149>, <http://dx.doi.org/10.1049/cit2.12149>, publisher: Institution of Engineering and Technology (IET)
48. Xing, J., Cioffi, G., Hidalgo-Carrió, J., Scaramuzza, D.: Autonomous Power Line Inspection with Drones via Perception-Aware MPC (2023), [_eprint: 2304.00959](https://arxiv.org/abs/2304.00959)
49. Ye, Y., Gandhi, D., Gupta, A., Tulsiani, S.: Object-centric forward modeling for model predictive control. In: Conference on Robot Learning. pp. 100–109. PMLR (2020)

50. Yoo, U., Zhao, H., Altamirano, A., Yuan, W., Feng, C.: Toward Zero-Shot Sim-to-Real Transfer Learning for Pneumatic Soft Robot 3D Proprioceptive Sensing. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 544–551 (2023). <https://doi.org/10.1109/ICRA48891.2023.10160384>
51. Zhang, K., Shi, Y., Sheng, H.: Robust Nonlinear Model Predictive Control Based Visual Servoing of Quadrotor UAVs. *IEEE/ASME Transactions on Mechatronics* **26**(2), 700–708 (2021). <https://doi.org/10.1109/TMECH.2021.3053267>
52. Zuo, Y., Qiu, W., Xie, L., Zhong, F., Wang, Y., Yuille, A.L.: CRAVES: Controlling Robotic Arm With a Vision-Based Economic System. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4209–4218 (2019). <https://doi.org/10.1109/CVPR.2019.00434>