Learning 2D Linear Dynamics in Image Space Using Deep Neural Networks

JEFFREY MAHLER, MICHAEL LASKEY, JOHN SCHULMAN, SERGEY LEVINE, SACHIN PATIL, PIETER ABBEEL, KEN GOLDBERG

RSS Workshop on Information Based Planning and Manipulation

Introduction

Grasping and manipulating a previously unknown object without assumptions about structure or relevant features in the environment is an especially challenging problem in robotics. Motivated by recent work using feature learning with deep neural networks, we approach this problem by learning a dynamical system that relates robotic control inputs (e.g., motor torques) to features learned from raw sensory data (e.g., images) using a deep neural network. We present an algorithm for inferring the most likely parameters of this system using approximate expectation maximization. Initial experiments show that the method is able to predict the next image of a pendulum in a 2D simulator to less than 1 pixel of mean squared error using only the current image and the control input to the system.

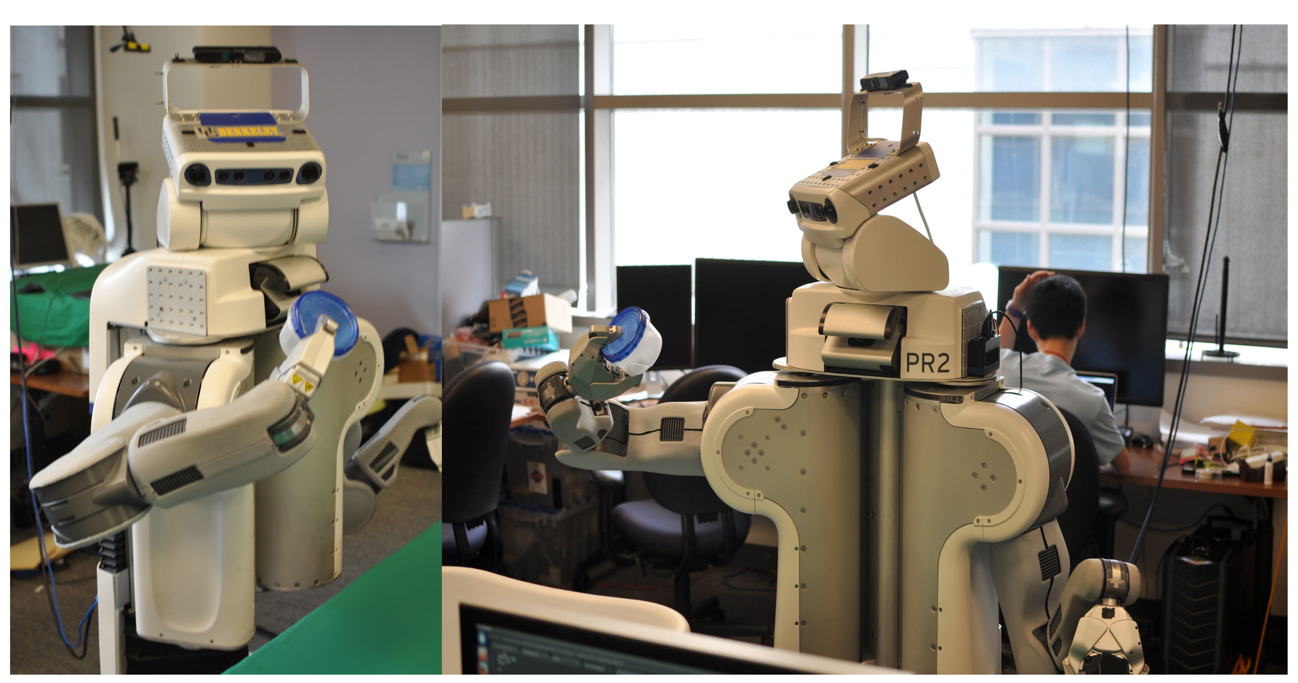

Motivation

- Visual Prediction of images can enable control of an object using strictly images
- The use of deep auto-encoders is motivated by their ability to learn a low dimensional state representation
- We try and learn a linear dynamics so methods like LQR can be used to control around a trajectory of images

Methods

- Learn a non-linear dynamical system from on a feature representation of images:

$$\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + \mathbf{v}_t \qquad \mathbf{v}_t \sim N(0, \sigma_1^2 I)$$
$$\mathbf{y}_t = h(\mathbf{x}_t) + \mathbf{w}_t \qquad \mathbf{w}_t \sim N(0, \sigma_2^2 I)$$

- EM Algorithm to learn dynamics
  - **E-Step:**
    - Kalman smoothing with EKF
  - **M-Step:**
    - Learn linear dynamics through weighted least squares
    - Stochastic Gradient Descent on samples from feature distribution to update deep net
- Observation function represented by Deep Neural Network
  - Maps hidden state to image space
  - Initialized by training an auto-encoder

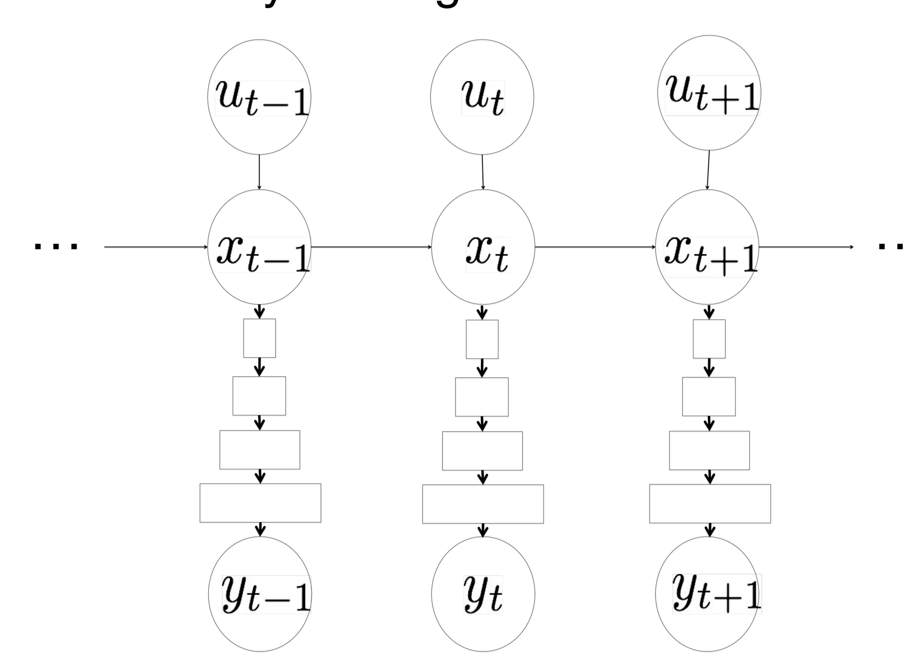

Future Work

- Extend work to natural images
- Implement a controller around learned dynamics and use it to reach desire images
- Combine state exploration approaches with our learning algorithm

Experiments

- Learned dynamics on pendulum in simulation
  - 9,000 training images and 1,000 test images with random controls
  - Trained for 1,000 iterations of EM
- Tried to predict future observations
  - Predicts several time steps ahead with low mean squared reconstruction error
  - Blurs out longer-term predictions due to uncertainty

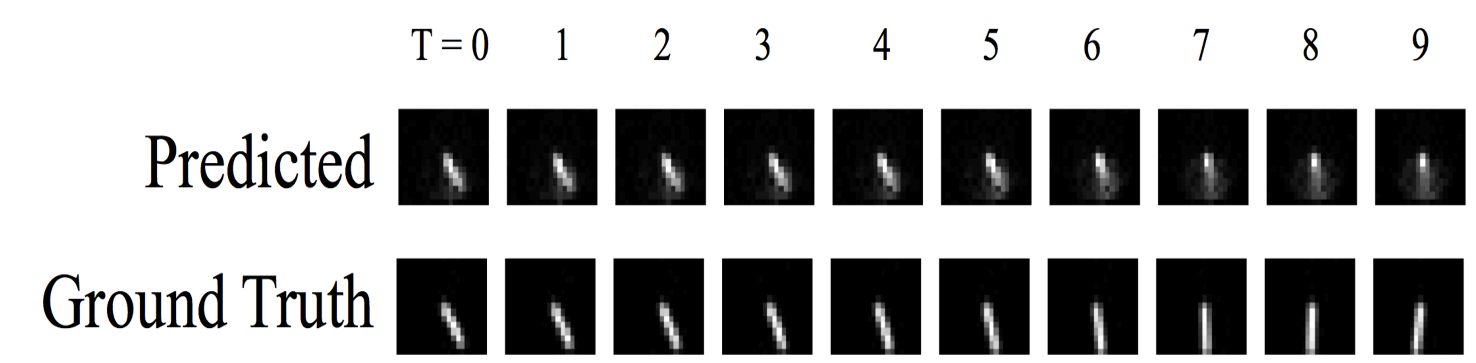

Fig. 1. Model predictions versus ground truth for a 10 image horizon

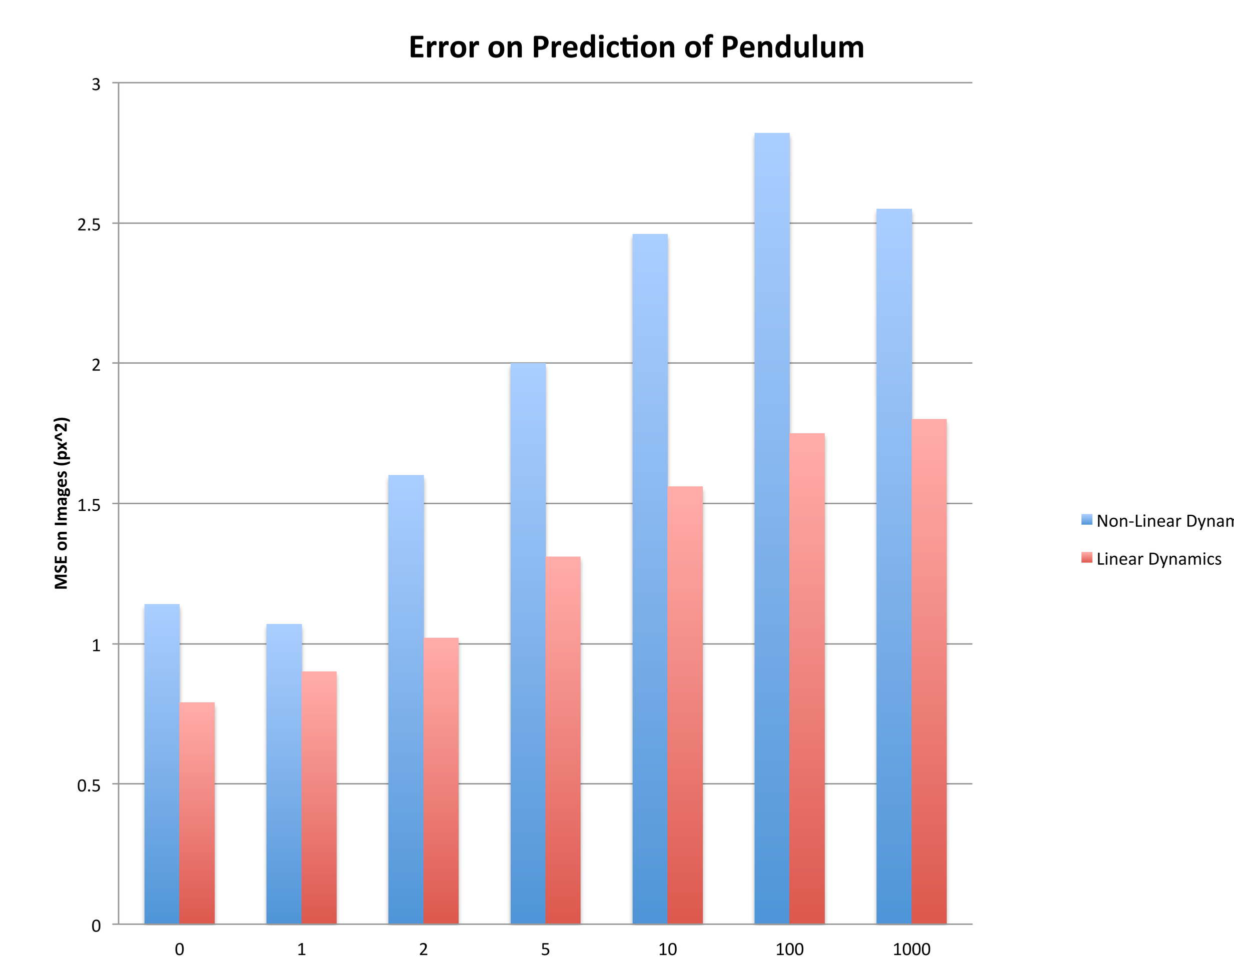

Fig 2. Horizon versus mean squared reconstruction error after 1,000 iterations of EM

Selected References

Hinton, G. E., Krizhevsky, A., & Wang, S. D. (2011). Transforming auto-encoders. In Artificial Neural Networks and Machine Learning–ICANN 2011 (pp. 44-51). Springer Berlin Heidelberg.

Weston, Jason, et al. "Deep learning via semi-supervised embedding." Neural Networks: Tricks of the Trade. Springer Berlin Heidelberg, 2012. 639-655.

Riedmiller, Martin. "Neural reinforcement learning to swing-up and balance a real pole." Systems, Man and Cybernetics, 2005 IEEE International Conference on. Vol. 4. IEEE, 2005.