Extracting Latent State Representations with Linear Dynamics from Rich Observations

Abraham Frandsen, Rong Ge, Holden Lee

Duke University

Control with Rich Observations

Challenges:

- irrelevant or redundant features
- complicated dynamics
- difficult to efficiently learn policies

Previous Work

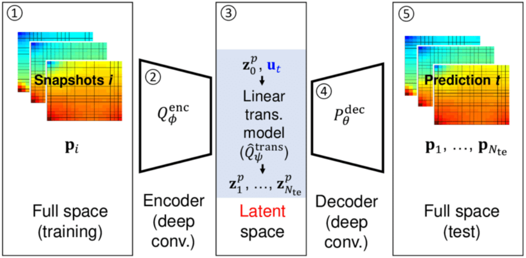

- Learn state representations with linear dynamics [WSBR15]
- Fit representations to forward or inverse dynamics models

Objectives of This Work

1. Analyze approaches based on forward and inverse models
2. Devise new algorithms with provable guarantees
3. Validate algorithms with simple experiments

Hidden Subspace Model

Observation x_t has a *linear* and *nonlinear* part:

$$x_t = y_t + z_t$$

Linear dynamics on a latent subspace V

$$\begin{aligned} &y_i \in V \\ &y_{t+1} = Ay_t + Bu_t \end{aligned}$$

Nonlinear, irrelevant dynamics on $V^\perp$:

$$\begin{aligned} &z_i \in V^\perp \\ &\text{corr}(z_{t+1}, (x_t, u_t)) < 1 \\ &z_{t+1} \perp (x_t, u_t) \,|\, y_{t+1} \end{aligned}$$

Forward Model

Predict next state given current state and action

Forward Model

Predict next state given current state and action

Goal: $Px_1 = APx_0 + Bu_0$

First Attempt:

$$\min_{P,A,B} \mathbb{E}\|Px_1 - (APx_0 + Bu_0)\|_2^2$$

Issues: trivial solutions, bad local minima

Forward Model

Regularized Objective:

$$\min_{P,Q,D} \frac{1}{2}\mathbb{E}\|Px_1 - Qx_0 - Du_0\|_2^2 + \frac{\lambda}{4}\|P\Sigma_{x_1}P^\top - I\|_F^2 \quad (1)$$

Forward Model

Regularized Objective:

$$\min_{P,Q,D} \frac{1}{2}\mathbb{E}\|Px_1 - Qx_0 - Du_0\|_2^2 + \frac{\lambda}{4}\|P\Sigma_{x_1}P^\top - I\|_F^2 \tag{1}$$

Theorem

Let (P^*, Q^*, D^*) be a local minimum of (1). Then the rows of P^* span the subspace V.

Inverse Model

Predict action given current and previous state

Inverse Model

Predict action given current and previous state

Observe: $u_0 = B^+ P x_1 - B^+ A P x_0$

$$\min_{P,A,C} \mathbb{E} \| u_0 - (CPx_1 - CAPx_0) \|_2^2$$

Issues: bad local minima, can't recover entire subspace V

Inverse Model

Multi-step convex relaxation:

$$\min_{\theta} \frac{1}{2}\mathbb{E}\sum_{i=1}^{r} \|u_{i-1} - (Px_i - L_i x_0 - \sum_{k=1}^{i-1} T_k u_{i-1-k})\|_2^2 \qquad (2)$$

Inverse Model

Multi-step convex relaxation:

$$\min_{\theta} \frac{1}{2}\mathbb{E}\sum_{i=1}^{r} \|u_{i-1} - (Px_i - L_i x_0 - \sum_{k=1}^{i-1} T_k u_{i-1-k})\|_2^2 \qquad (2)$$

Theorem

Let $P^*, L_1^*, \cdots, T_1^*, \cdots$ be a minimal-norm optimal solution of (2). Then the rows of $P^*, L_1^*, \ldots, L_r^*$ span the subspace V.

Experiment: Inverted Pendulum

Task: swing the pendulum into vertical position, stabilize

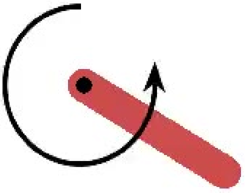

Policy Learning

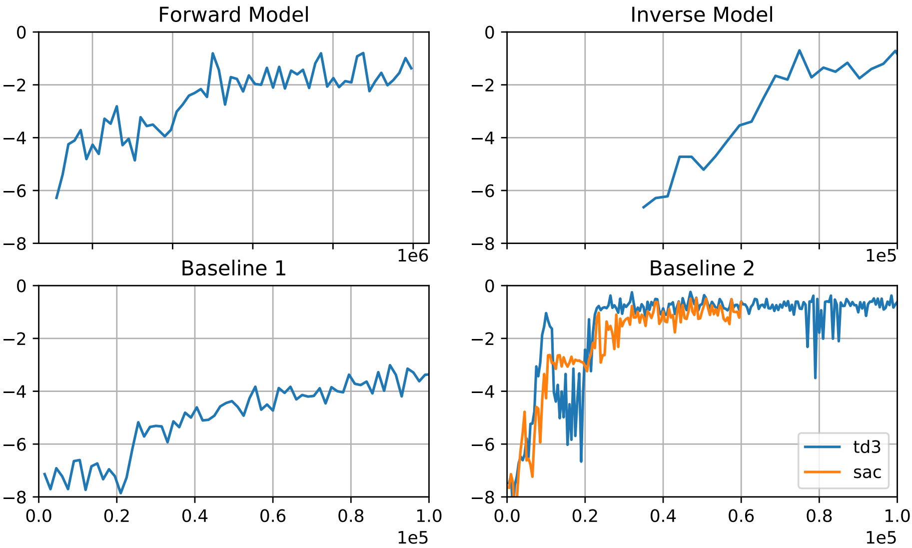

References

Manuel Watter, Jost Springenberg, Joschka Boedecker, and Martin Riedmiller, *Embed to control: A locally linear latent dynamics model for control from raw images*, Advances in neural information processing systems, 2015, pp. 2746–2754.