Maximum-Likelihood Learning of Latent Dynamics Without Reconstruction

Samo Hromadka, Kai Biegun, Lior Fox, James Heald, Maneesh Sahani

Gatsby Unit, University College London

GATSBY

UCL

Summary

- We extend the recognition-parametrized model (RPM[1]) to latent time series.
- Unlike *generative* methods, the RPM does not learn an explicit generative model.
  - $\Rightarrow$ Does not spend capacity on reconstructing distractors.
- Unlike *contrastive* methods, the RPM computes posterior uncertainty in a maximum-likelihood framework.
  - $\Rightarrow$ Less prone to model collapse.
- We apply our method to distractor-heavy videos of physical dynamical systems, object tracking, and continuous control tasks.

Setup

- We assume a probabilistic state space model and consider tasks where downstream inference (recognition) is more important than generation.

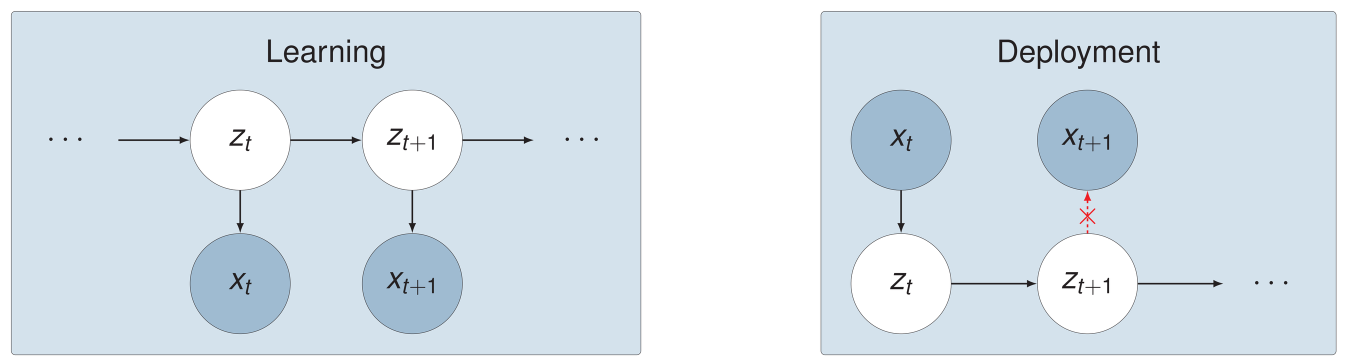

$$p(x_{1:T}, z_{1:T}) = p(z_1)\prod_{t=2}^{T} p(z_t|z_{t-1})\prod_{t=1}^{T} p(x_t|z_t) \tag{1}$$

Related Work

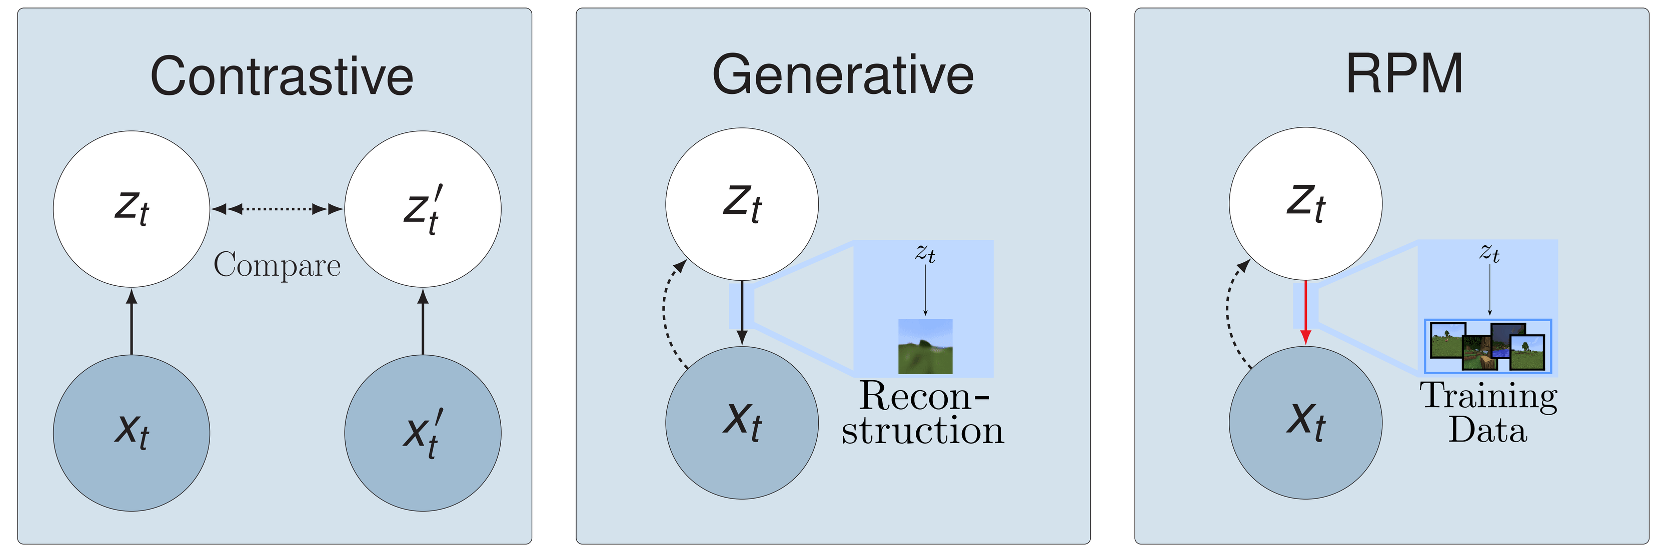

The Model

- The **recognition-parametrized Gaussian state space model** (RP-GSSM) reparametrizes equation (1) as

$$p_\theta^{\boldsymbol{X}}(x_{1:T}, z_{1:T}) = p_\eta(z_1)\prod_{t=2}^{T} p_\eta(z_t|z_{t-1})\prod_{t=1}^{T}\frac{f_{\phi_t}(z_t|x_t)p_0^{\boldsymbol{X}}(x_t)}{F_{\phi_t}(z_t)}, \tag{2}$$

where:

- $\boldsymbol{X}$ is a dataset of N time series observations: $\boldsymbol{X} = (x_1^n, \ldots, x_T^n)_{n=1}^N$;
- p_η is a parametrization of the prior (linear-Gaussian);
- $f_{\phi_t}(z_t|x_t)$ are parametrized distributions over z_t (nonlinear-Gaussian);
- $p_0^{\boldsymbol{X}}$ is the empirical data distribution: $p_0^{\boldsymbol{X}}(x_t) = \frac{1}{N}\sum_{n=1}^N \delta(x_t = x_t^n)$;
- $F_{\phi_t}(z_t)$ ensures that the joint distribution is normalized:

$$F_{\phi_t}(z_t) = \int f_{\phi_t}(z_t|x_t)p_0^{\boldsymbol{X}}(x_t)\,\mathrm{d}x_t = \frac{1}{N}\sum_{n=1}^{N} f_{\phi_t}(z_t|x_t^n).$$

- Learning is done by maximizing the free energy (ELBO)

$$\mathcal{F}(q,\theta) = \mathbb{E}_{q(z_{1:T})}\left[\log p_\theta^{\boldsymbol{X}}(x_{1:T}\,|\,z_{1:T})\right] - \mathrm{KL}(q(z_{1:T})\,\|\,p_\eta(z_{1:T}))$$

- For fixed θ, the optimal q can be found via Kalman smoothing.

Theorem (Consistency)

Let $\tilde{p}_\theta$ be a *limiting* RP-GSSM given by eq. (2) but with the true data distribution $p(x_t)$. Let $\boldsymbol{X}$ be a dataset of N i.i.d. data sequences sampled from $\tilde{p}_{\theta^*}$. Let $\hat{\mathcal{L}}_N(\theta)$ be the data likelihood under eq. (2), and let $\hat{\theta}_N$ be its maximizer. Then, under mild regularity conditions,

$$\hat{\theta}_N \xrightarrow[N\to\infty]{\mathbb{P}} \theta^*.$$

Theorem (Linear Identifiability)

Let $\tilde{p}_\theta$ and $\tilde{p}_{\theta'}$ be two limiting RP-GSSMs. Under mild regularity conditions, if $\tilde{p}_\theta(x_{1:T}) = \tilde{p}_{\theta'}(x_{1:T})$, then the posteriors $\tilde{p}_\theta(z_{1:T}|x_{1:T})$ and $\tilde{p}_{\theta'}(z_{1:T}|x_{1:T})$ are identical up to an invertible linear transformation.

Experiments: Uncontrolled Dynamics

- We train on videos with and without background distractors.
- Performance is measured by downstream linear regression R^2 to true states.

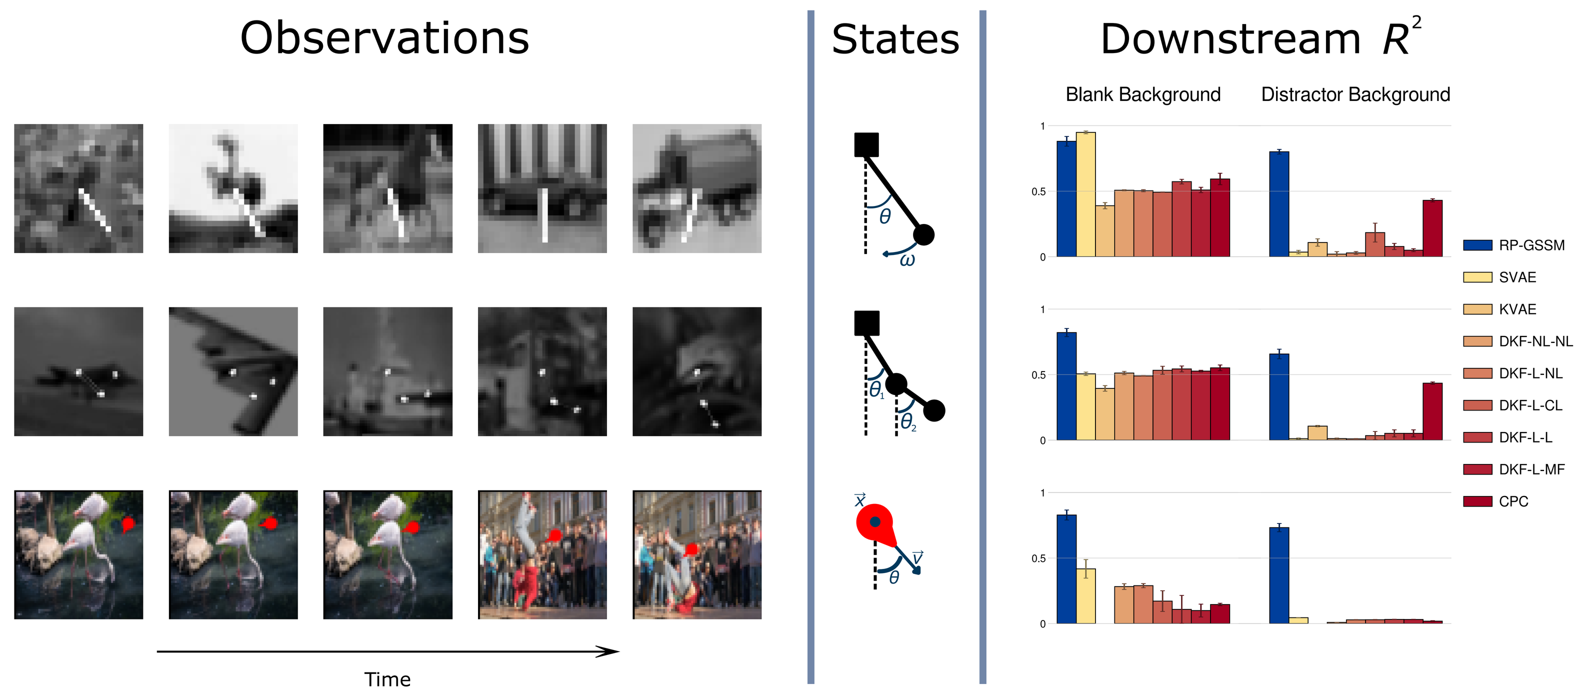

- To visualize learned features, we train an *auxiliary generative model* on top of each model's frozen representations.

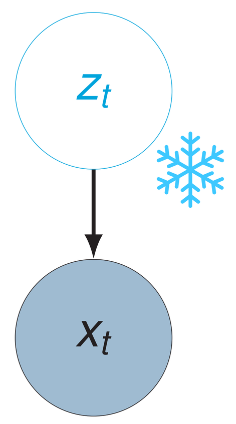

- The auxiliary model is trained to minimize the usual generative loss:

$$-\mathbb{E}_{p(x)}\mathbb{E}_{q(z)}\left[\log p_\psi(x|z)\right]$$

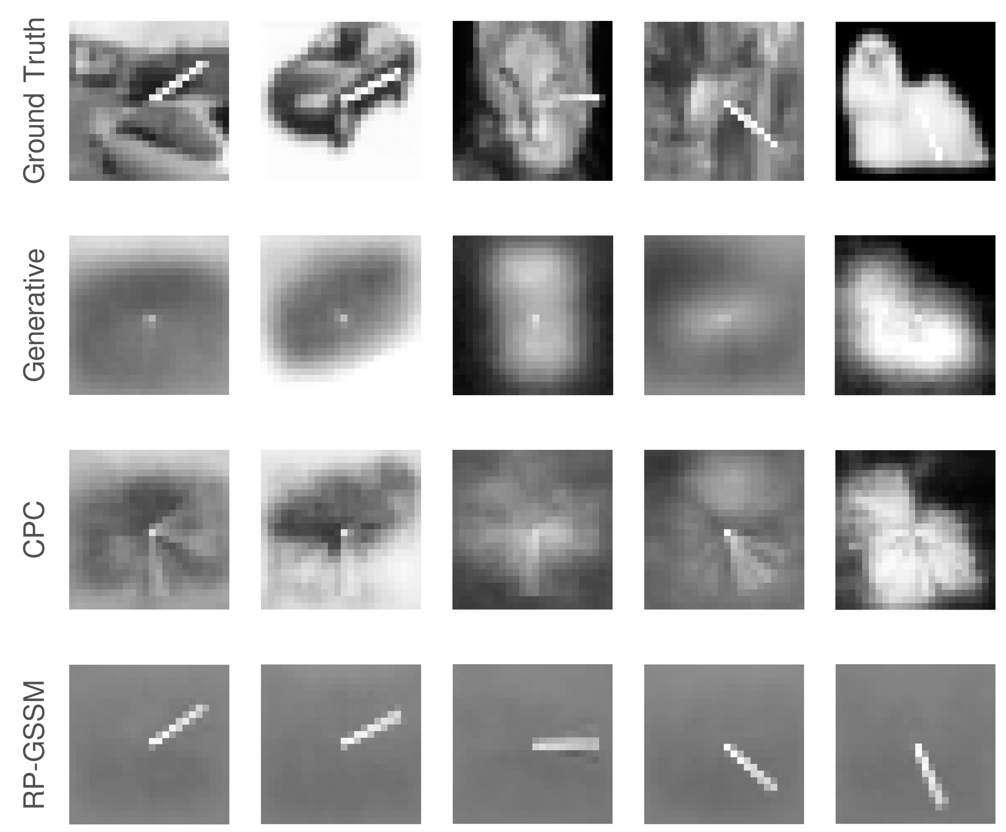

Experiments: Controlled Dynamics

- The RP-GSSM naturally generalizes to dynamics with control inputs.
- Inference can be performed exactly even when control influence is nonlinear:

$$p_\eta(z_{t+1}|z_t, a_t) = \mathcal{N}\left(A_\eta(a_t)z_t + b_\eta(a_t), Q_\eta(a_t)\right).$$

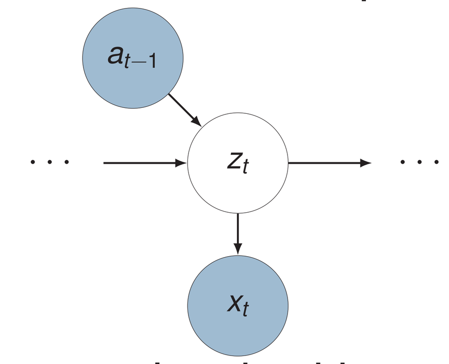

- We test performance on the Walker continuous control task with background videos and camera jitter.
- Kernel ridge regression R^2 to true joint angles, velocities, and height:

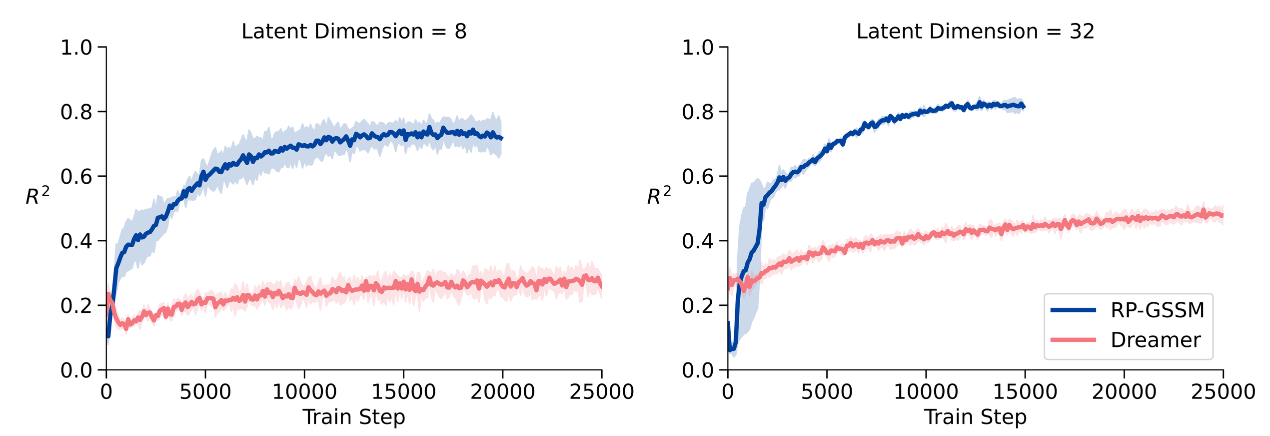

- Auxiliary reconstructions of model rollouts:

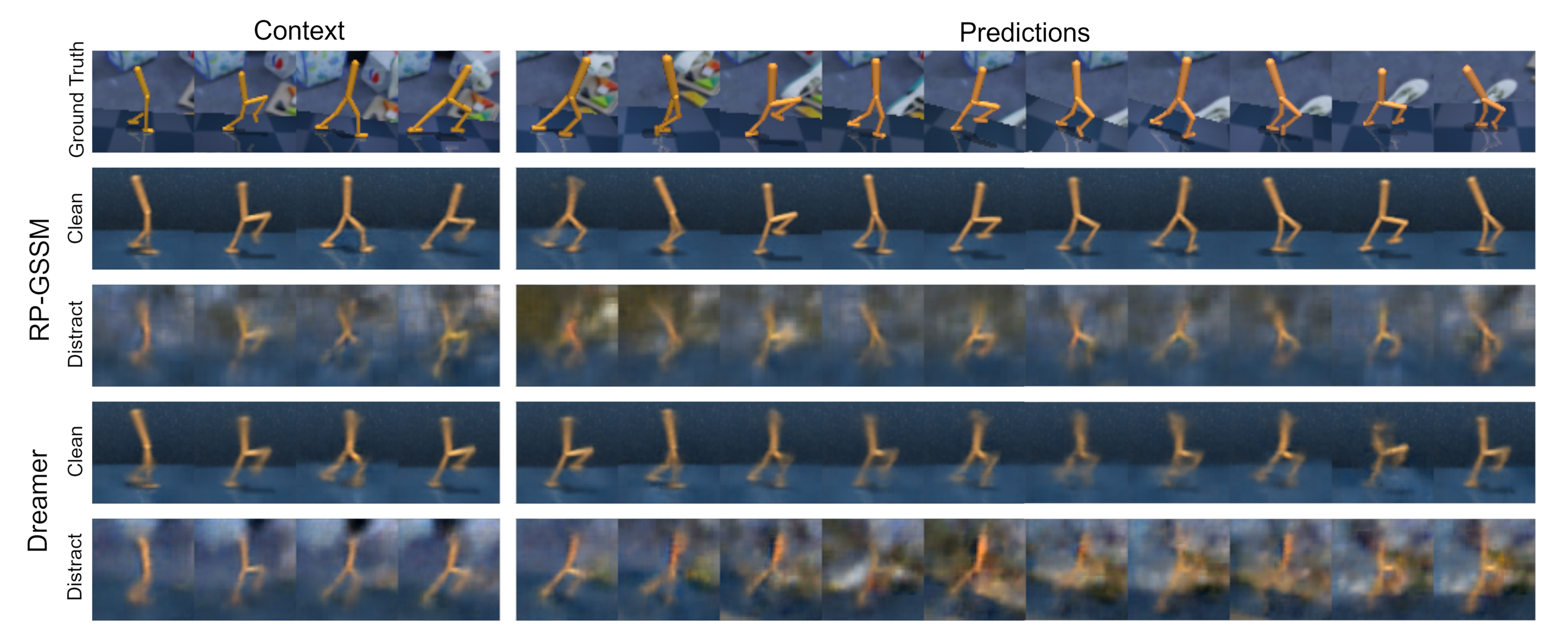

References

[1] William I. Walker, Hugo Soulat, Changmin Yu, and Maneesh Sahani. Unsupervised representation learning with recognition-parametrised probabilistic models. In *AISTATS*, 2023.