北京理工大學
BEIJING INSTITUTE OF TECHNOLOGY

NEURAL INFORMATION PROCESSING SYSTEMS

DyMoDreamer: World Modeling with Dynamic Modulation

Advances in Neural Information Processing Systems (NeurIPS 2025)

Boxuan Zhang, Runqing Wang, Wei Xiao, Weipu Zhang, Jian Sun, Gao Huang, Jie Chen, Gang Wang

School of Automation, Beijing Institute of Technology
Department of Automation, BNRist, Tsinghua University

➢ World model-based reinforcement learning

- Learning latent dynamics for planning and decision-making;
- Generating ``imagination trajectories” for agents to learn behaviors.

Environment Interaction

Learned World Model

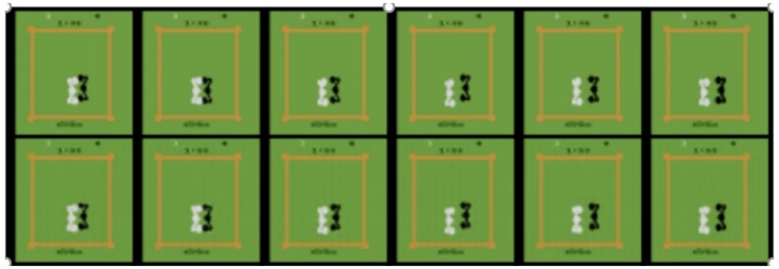

Planning by Latent Imagination

Motivation

➢ Cognitive Science

- Human infants naturally focus on dynamic object interactions to infer fundamental principles of their surroundings;
- Rewards are predominantly influenced by the dynamic parts.

Atari 100k – Pong/Krull

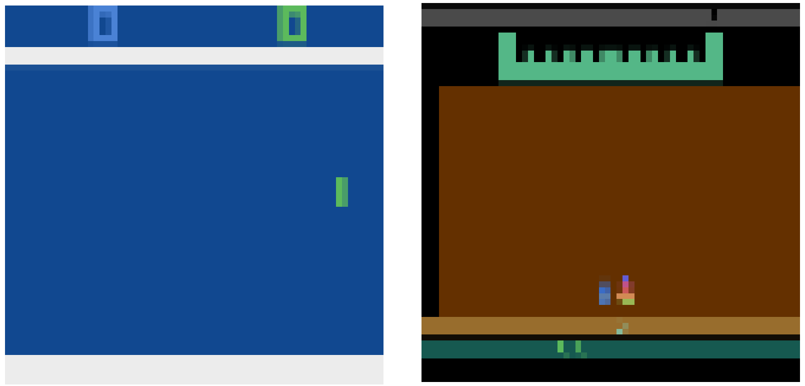

- **Conventional world models process observations holistically;**

- **Designing a world model focus on the key dynamic features and temporal information.**

Method

➢ DyMoDreamer

- Prioritizing dynamic features through inter-frame differencing mask;
- Integrating a dynamic modulation mechanism into RSSM;
- All components are trained jointly without separate modulator training.

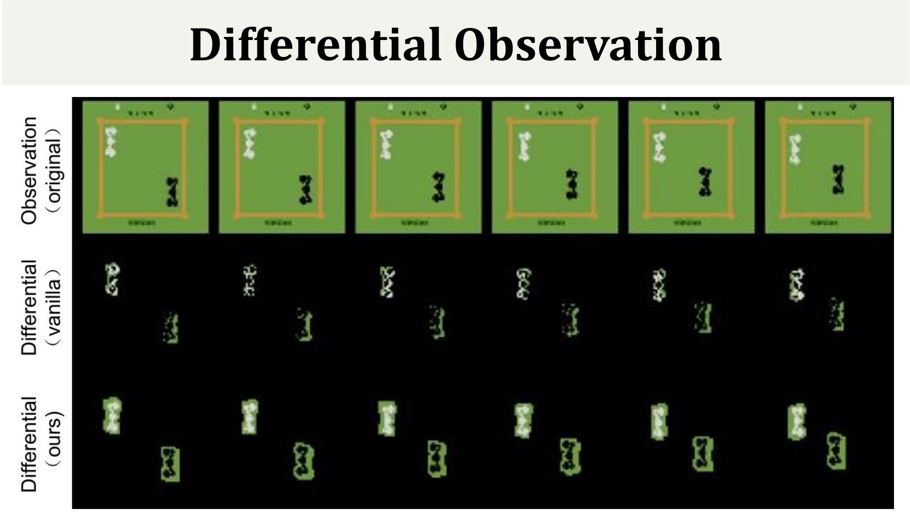

Method

➢ DyMoDreamer

- Prioritizing dynamic features through inter-frame differencing mask;
- Integrating a dynamic modulation mechanism into RSSM;
- All components are trained jointly without separate modulator training.

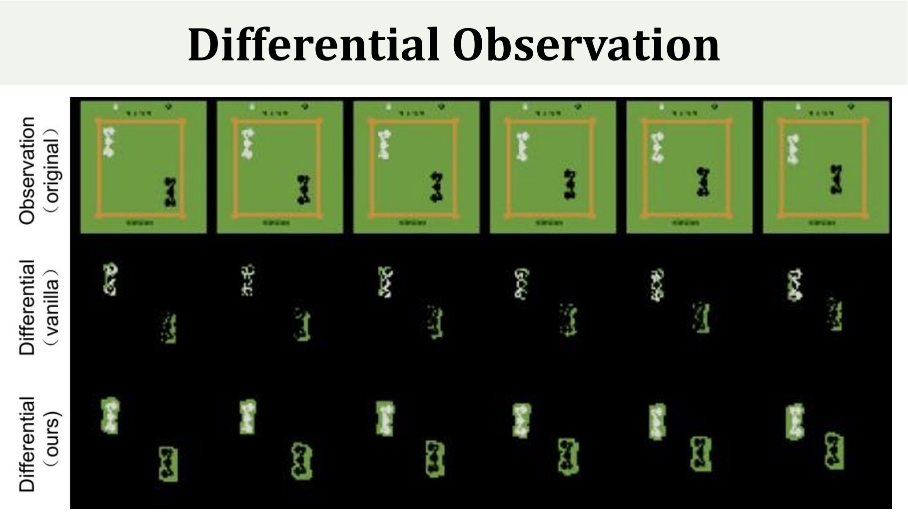

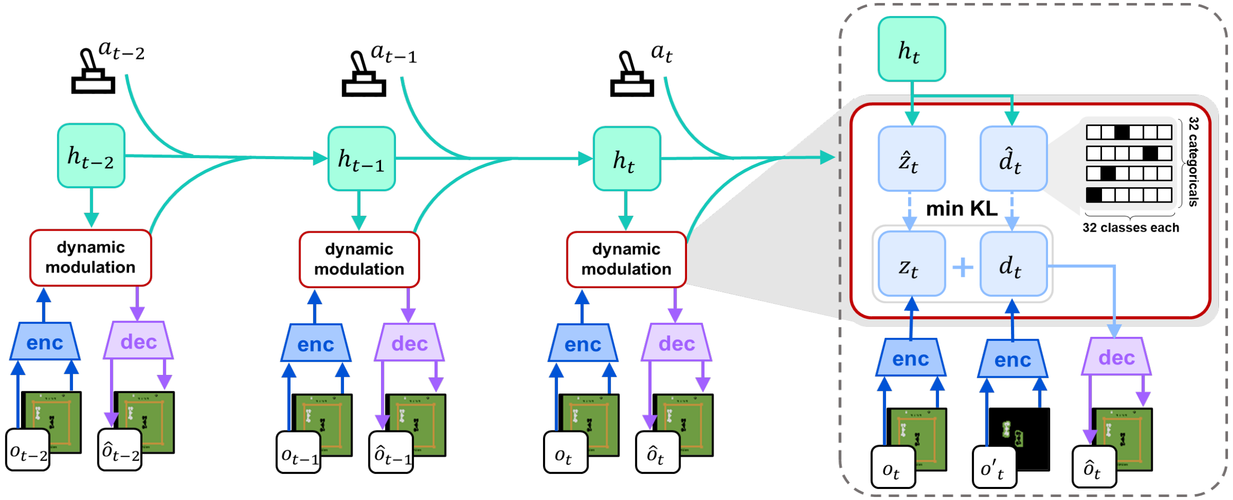

Method

➢ DyMoDreamer

- Prioritizing dynamic features through inter-frame differencing mask;
- Integrating a dynamic modulation mechanism into RSSM;
- All components are trained jointly without separate modulator training.

Prediction Loss

$$\mathcal{L}_{\text{pred}}(\phi) = \mathcal{L}_{\text{rec}}(\phi) + \mathcal{L}_{\text{rew}}(\phi) + \mathcal{L}_{\text{con}}(\phi)$$

$$\begin{aligned}
\mathcal{L}_{\text{rec}}(\phi) &= -\ln p_\phi(o_t|h_t, z_t, d_t) \\
\mathcal{L}_{\text{rew}}(\phi) &= -\ln p_\phi(r_t|h_t, z_t, d_t) \\
\mathcal{L}_{\text{con}}(\phi) &= -\ln p_\phi(c_t|h_t, z_t, d_t)
\end{aligned}$$

Dynamic and Representation Loss

$$\begin{aligned}
\mathcal{L}_{\text{dyn}}(\phi) = \max\left(1, \text{KL}[\text{sg}(q_\phi(z_t \mid h_t, o_t)) \parallel p_\phi(\hat{z}_t \mid h_t)]\right) \\
+ \underbrace{\max(1, \text{KL}[\text{sg}(q_\phi(d_t \mid h_t, o'_t)) \parallel p_\phi(\hat{d}_t \mid h_t)])}_{\text{modulation's dynamic loss}}.
\end{aligned}$$

$$\begin{aligned}
\mathcal{L}_{\text{rep}}(\phi) = \max\left(1, \text{KL}[q_\phi(z_t \mid h_t, o_t) \parallel \text{sg}(p_\phi(\hat{z}_t \mid h_t))]\right) \\
+ \underbrace{\max(1, \text{KL}[q_\phi(d_t \mid h_t, o'_t) \parallel \text{sg}(p_\phi(\hat{d}_t \mid h_t))])}_{\text{modulation's representation loss}}.
\end{aligned}$$

Divergence Regularization

$$\mathcal{L}_{\text{reg}}(\phi) = \frac{1}{B}\sum \text{KL}(\sigma(\Delta\hat{o}) \parallel \sigma(\Delta o)).$$

$$\sigma(\Delta\hat{o}_t) = \frac{\exp(\Delta\hat{o}_t/\tau)}{\sum_{h,w,c}\exp(\Delta\hat{o}_t/\tau)},$$

$$\sigma(\Delta o_t) = \frac{\exp(\Delta o_t/\tau)}{\sum_{h,w,c}\exp(\Delta o_t/\tau)},$$

➢ More Efficient Imagination

- DyMoDreamer highlights the dynamic features, exhibits less hallucinations during the imagination reconstruction than vanilla DreamerV3.

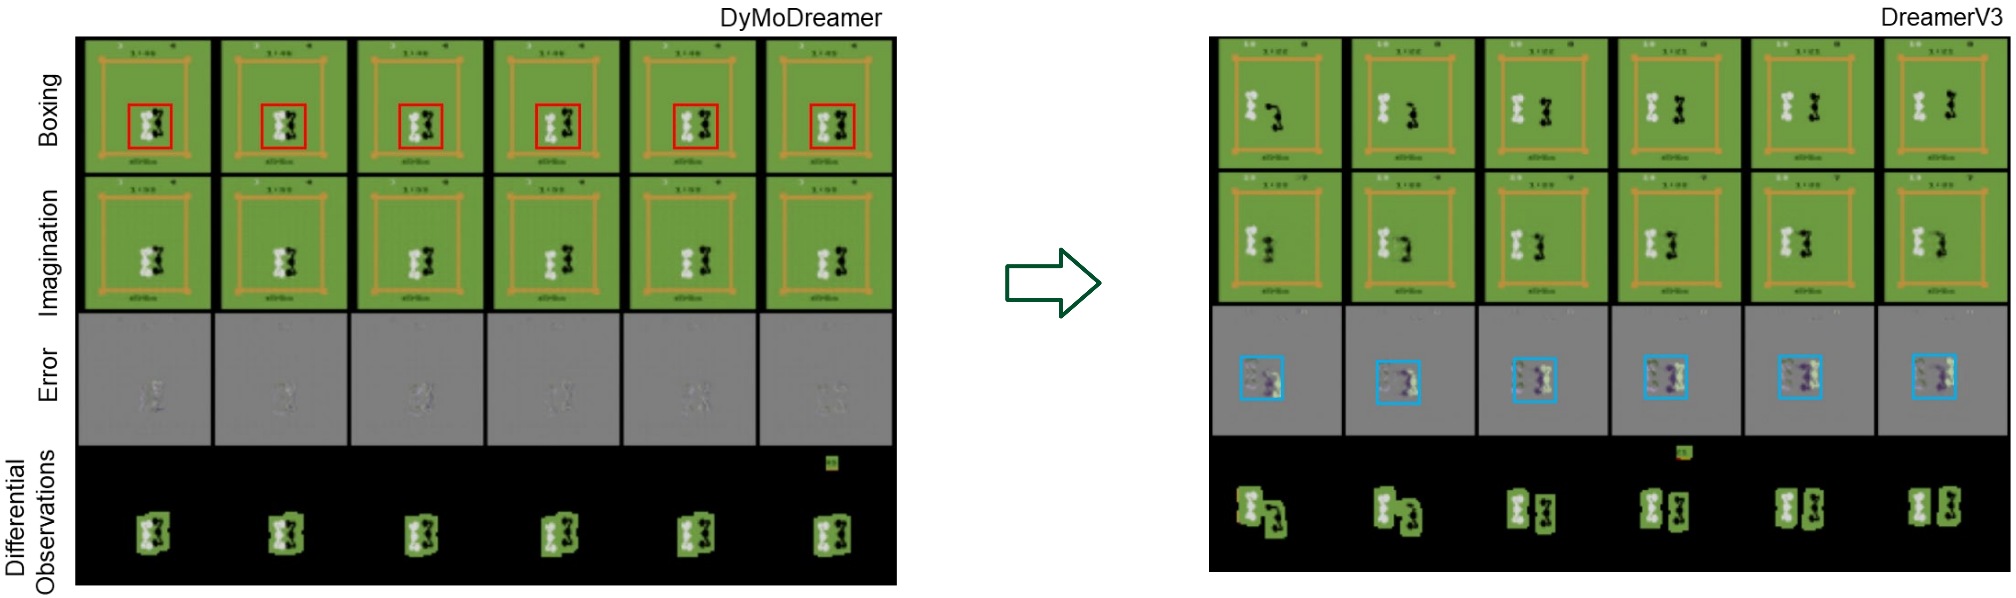

➢ Performance

- DyMoDreamer surpasses baselines on the Atari 100K, DeepMind Visual Control, and Crafter benchmark.

Atari 100k				
	OC-STORM (2025)	DIAMOND (2024)	DreamerV3 (2023)	DyMoDreamer (ours)
HNS Mean	134.8%	146%	125%	**156.6%**
HNS Median	43.8%	37%	49%	**71.3%**
DeepMind Visual Control Suite				
	TD-MPC2 (2023)	TWISTER (2025)	DreamerV3 (2023)	DyMoDreamer (ours)
Task Mean	720.9	801.8	786	**832**
Task Median	795.9	**907.6**	861	871
Crafter				
	IRIS (2023)	Δ-IRIS (2024)	DreamerV3 (2023)	DyMoDreamer (ours)
Return @1M	5.5	7.7	9.4	**10.3**

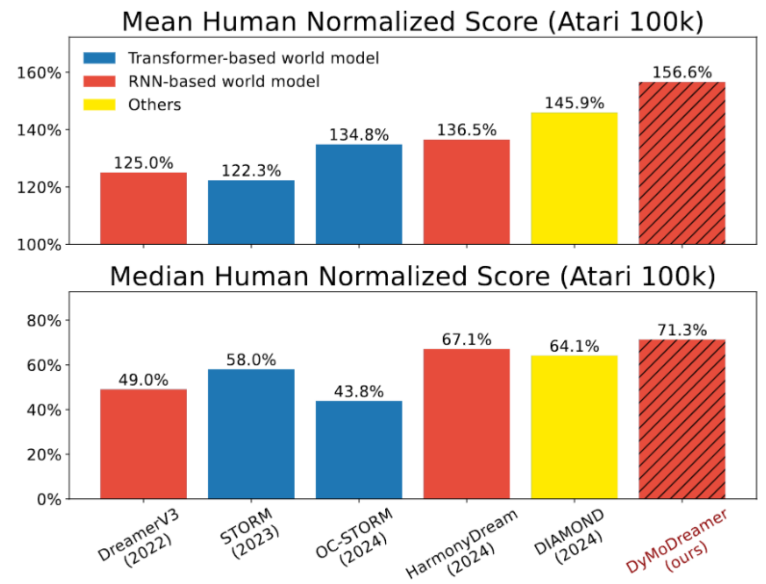

Ablation Studies

➢ Removing Dynamic Modulation

- Simply appending differential observations fails to direct the world model's focus toward dynamic patterns.

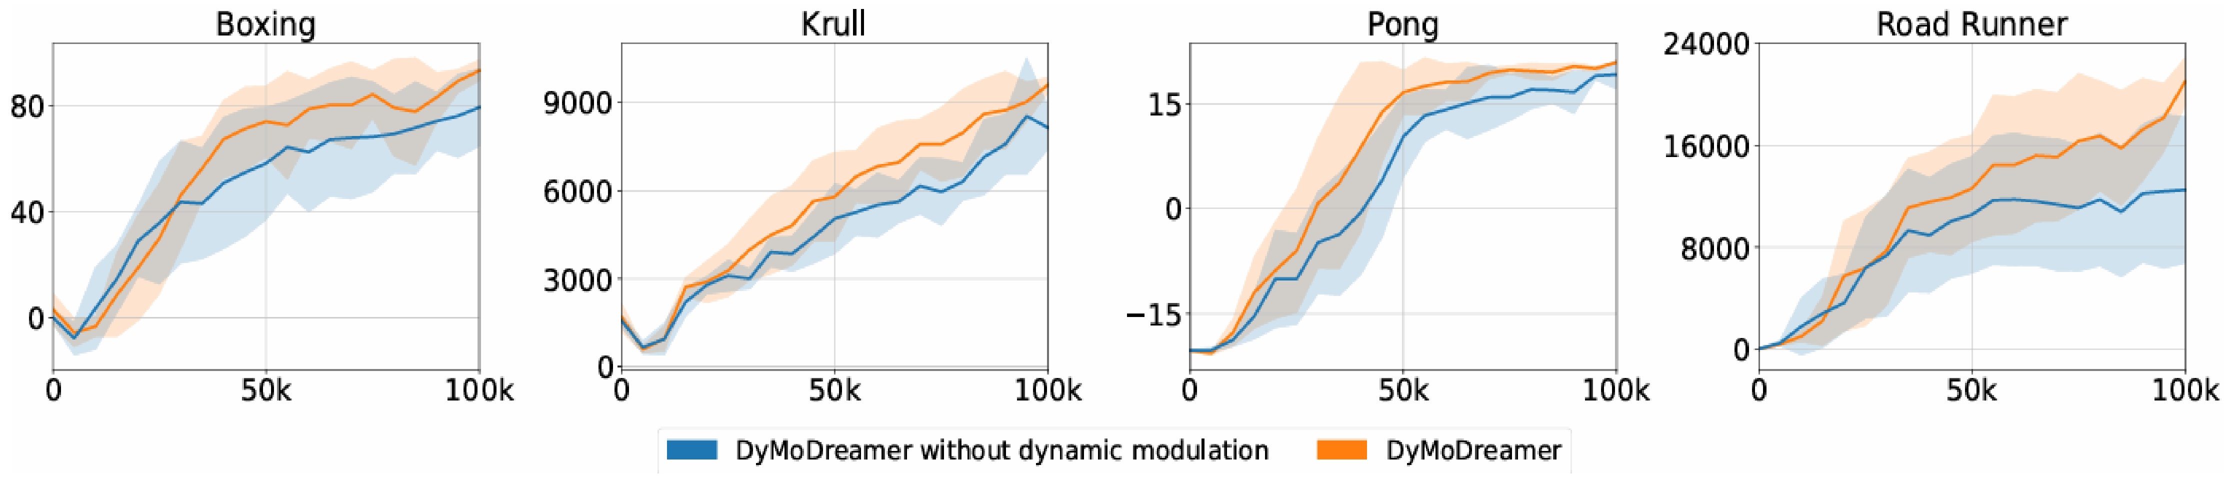

北京理工大學
BEIJING INSTITUTE OF TECHNOLOGY

DyMoDreamer: World Modeling with Dynamic Modulation

Advances in Neural Information Processing Systems (NeurIPS 2025)

Boxuan Zhang, Runqing Wang, Wei Xiao, Weipu Zhang,
Jian Sun, Gao Huang, Jie Chen, Gang Wang