

Task-Relevant Reconstruction for Visual Control with Distraction

Kyungmin Kim, Charless Fowlkes, and Roy Fox
University of California, Irvine

1. Visual Control with Distraction

- ❖ **Visual control task:** Control actions based on visual information.
 - e.g. DeepMind Control suite (DMC)
- ❖ Add **distractions** for a more **challenging** and **realistic** setup.



2. Model-Based Reinforcement Learning (MBRL)

- ❖ **Cooperation** between a **world model** and **behavior learning**.
- ❖ Promising with **great sample efficiency** in visual control tasks.
- ❖ Often **struggles** in **distracting** environments.

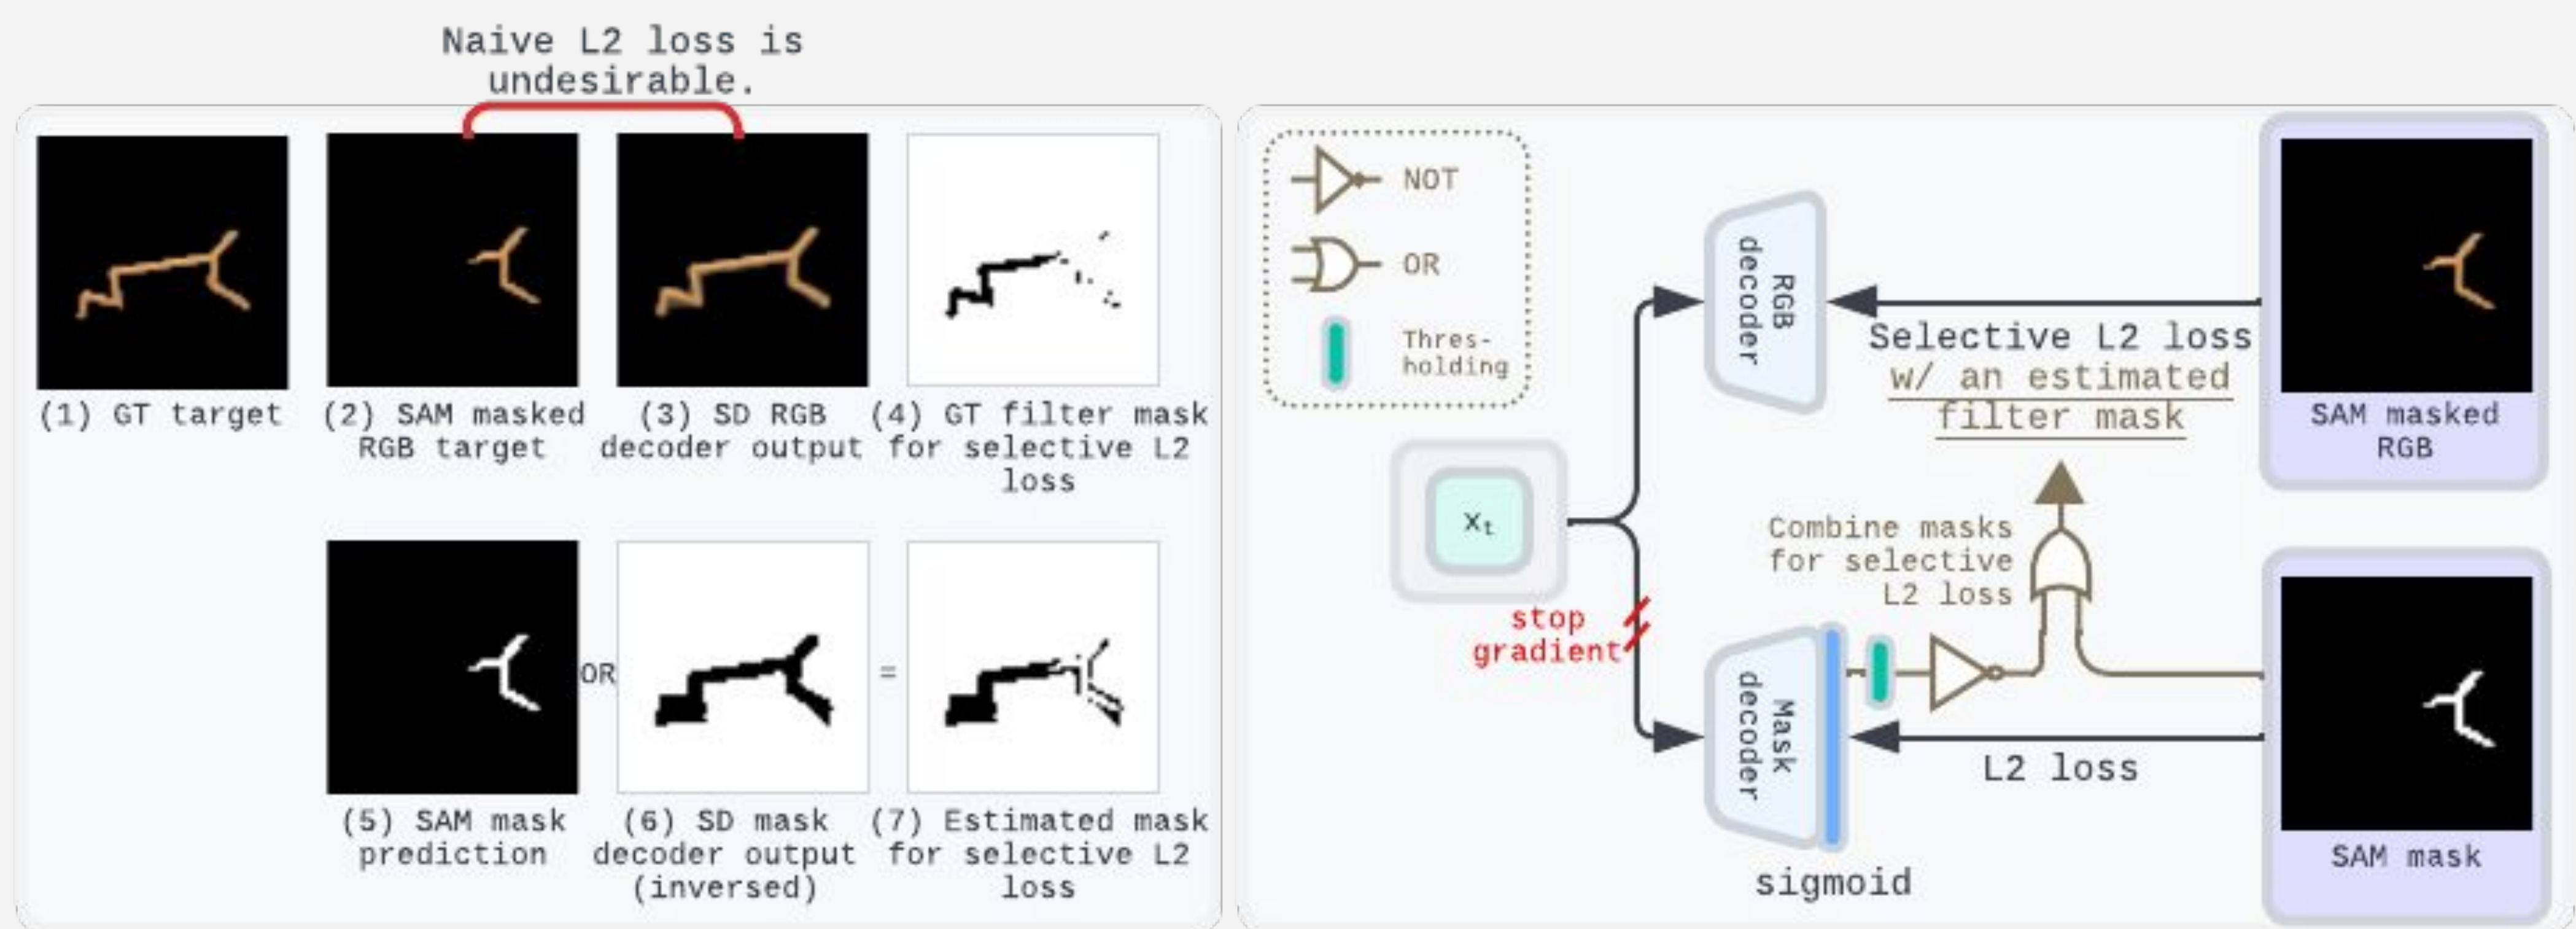
Representation learning	Examples	Drawbacks
Reconstruction-based	Dreamer [1], etc.	Unnecessary information included
Reconstruction-free	TD-MPC [2], DreamerPro [3], etc.	Sample inefficient

3. Method

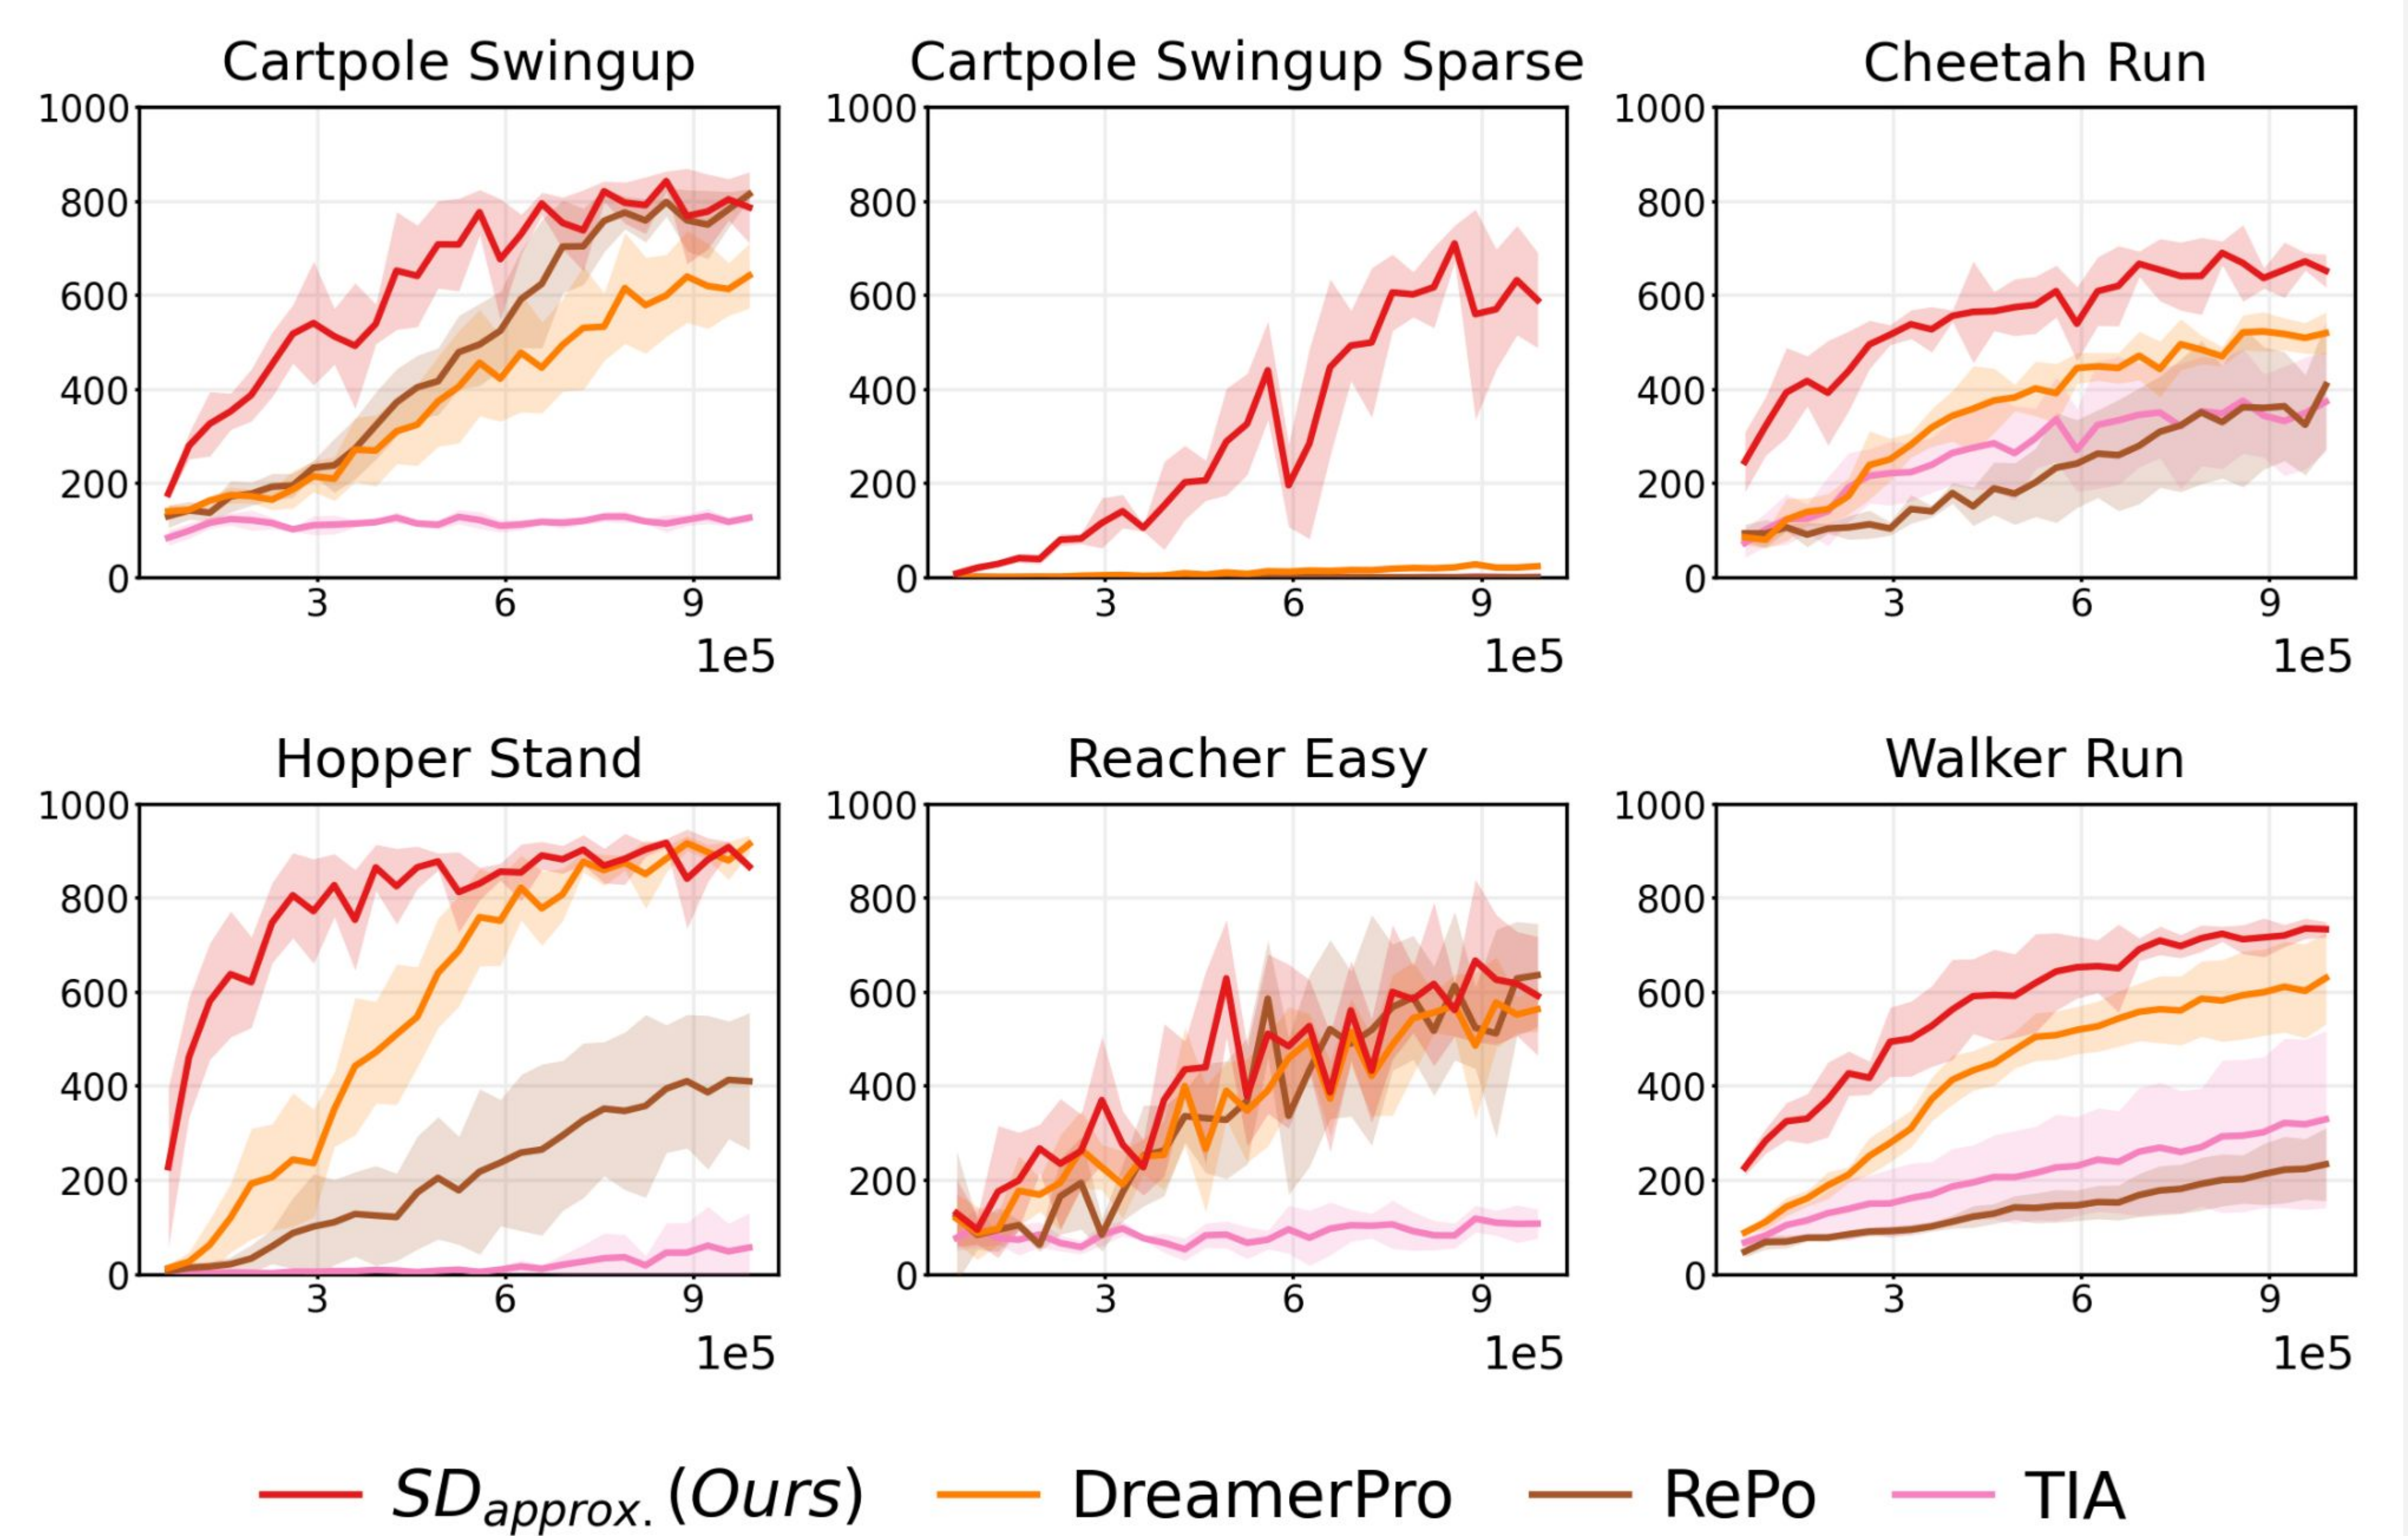
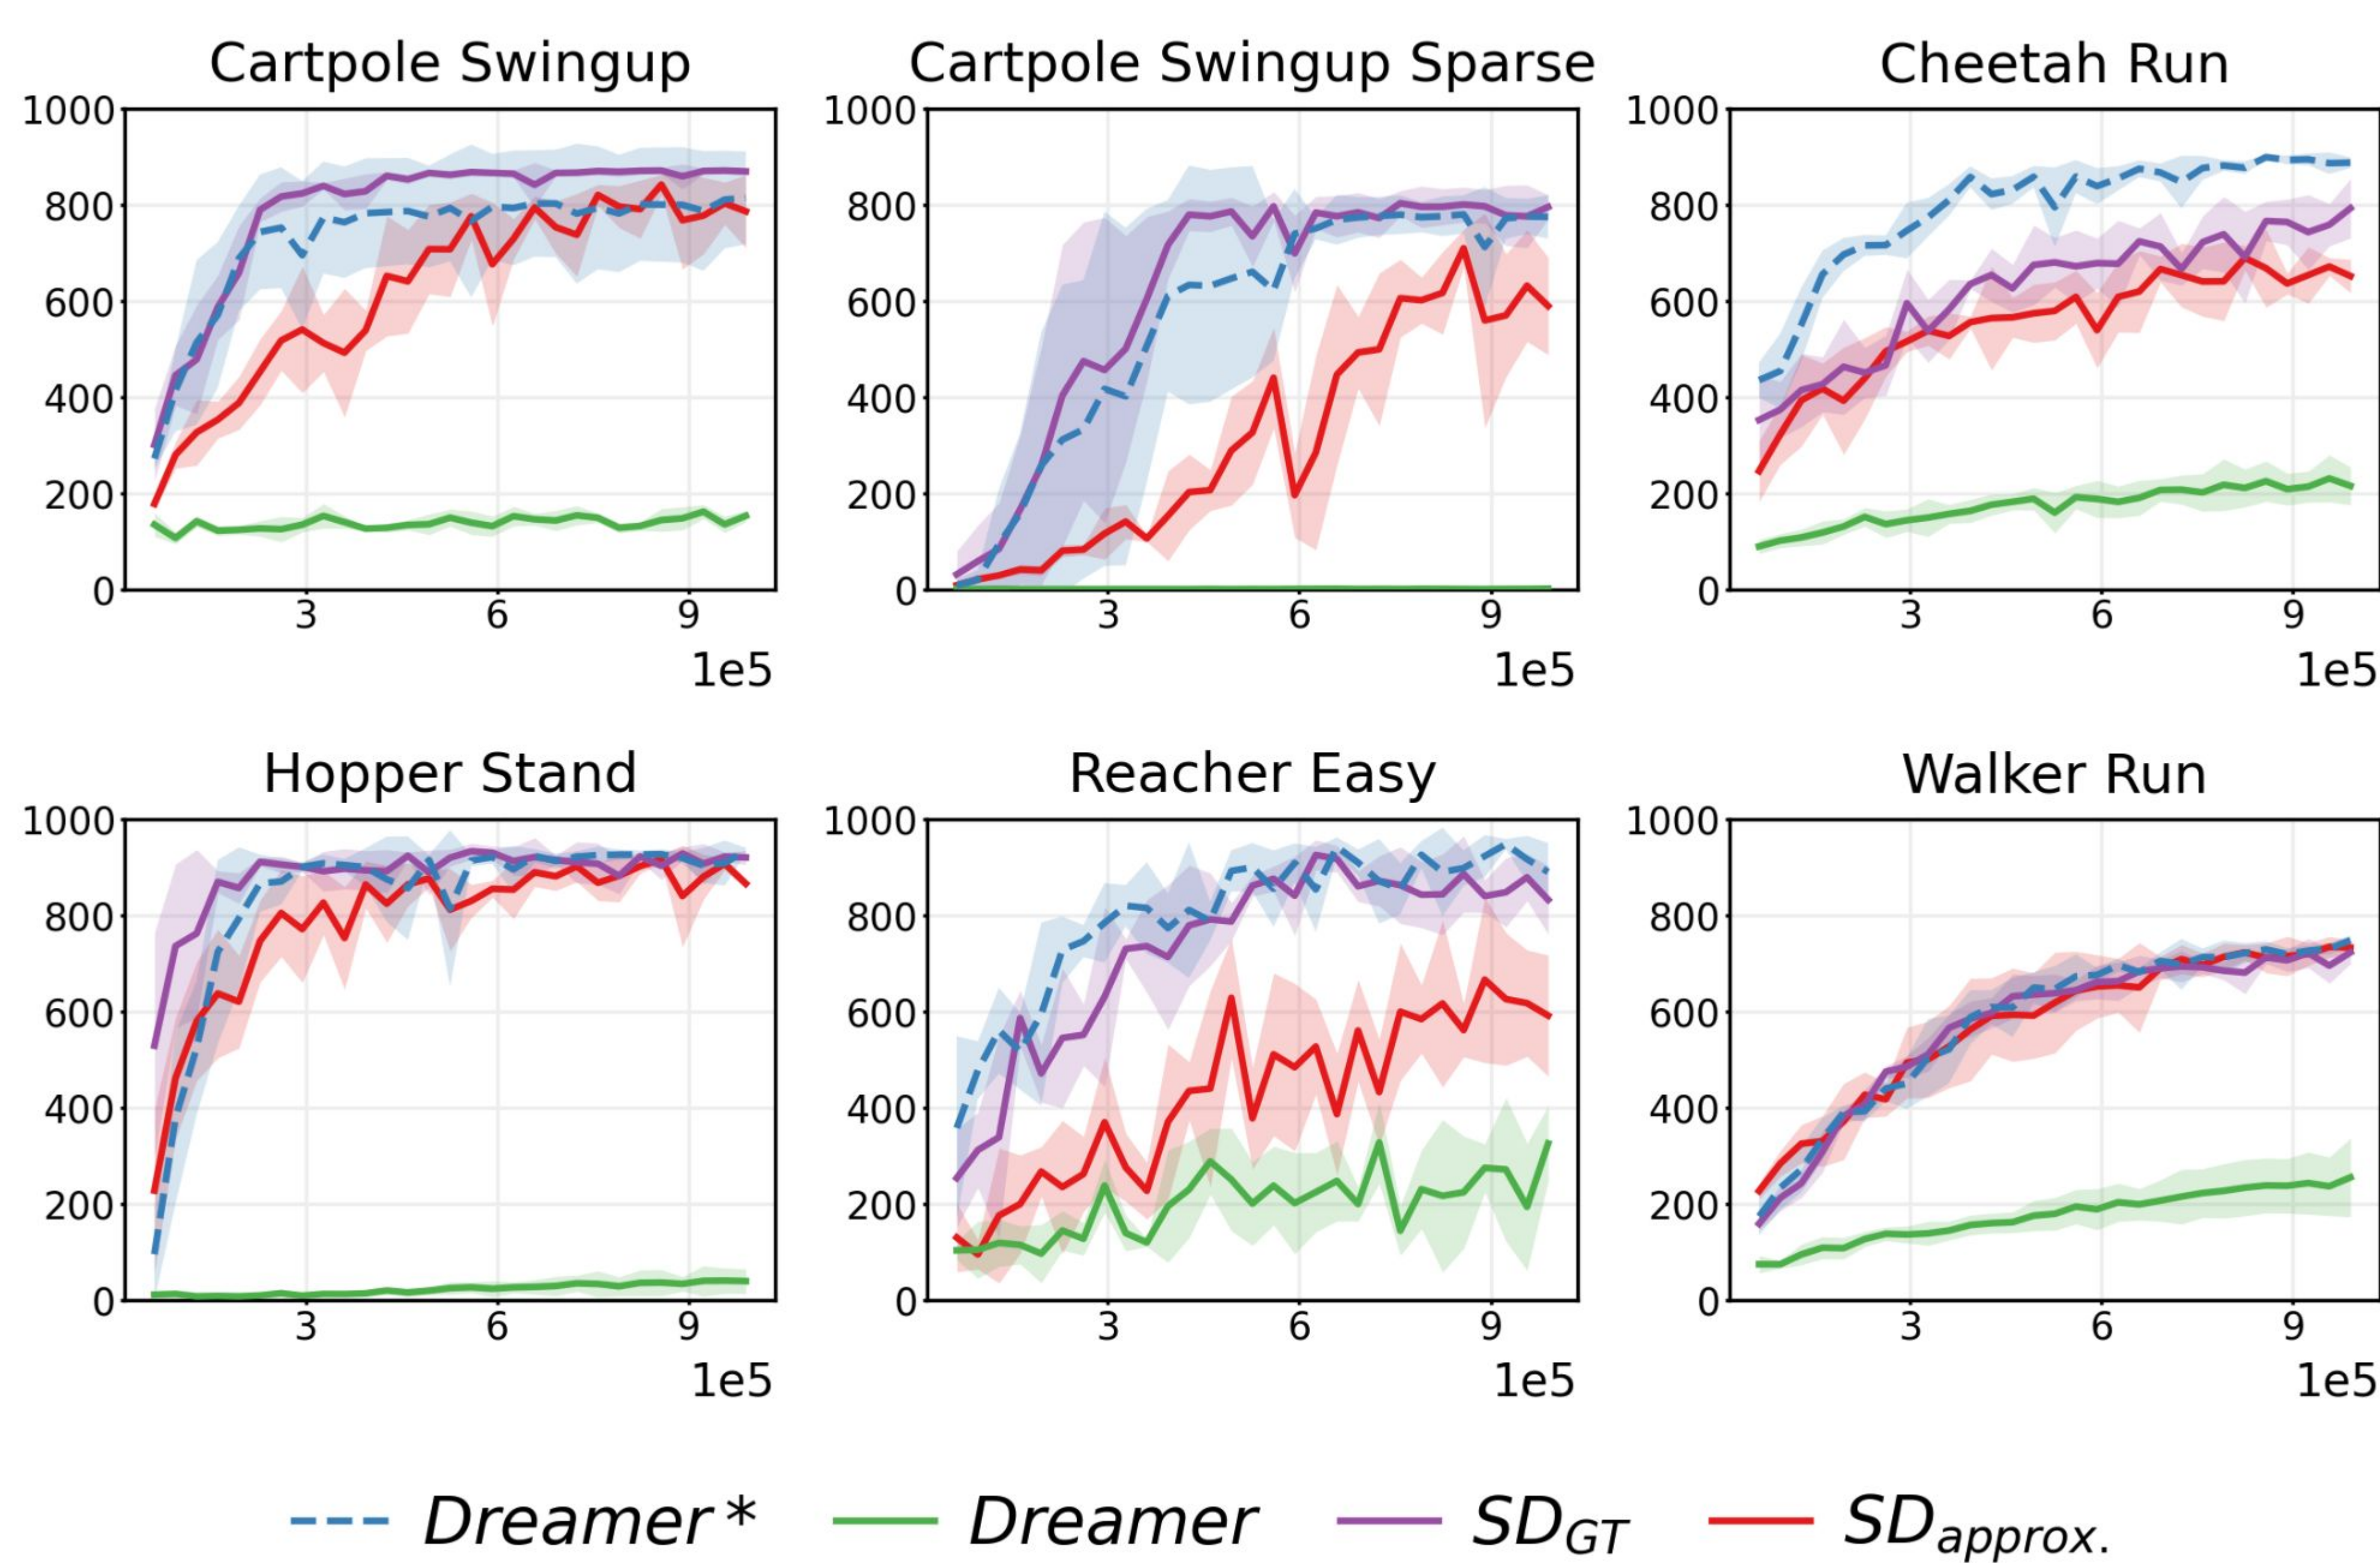
- ❖ **Assumption:** Task-relevant components are **straightforward to identify** within images, given available prior knowledge.
- ❖ Use **prior knowledge** with **segmentation foundation models**.
- ❖ **SD:** Reconstruct **task-relevant components only**.
- ❖ **SD_{GT}:** When **GT mask** for task-relevant components are **available**.
- ❖ **SD_{approx.}:** Use **PerSAM** [4] fine-tuned with a **single data point**.

- ❖ To make **SD_{approx.}** more **robust to noisy targets**, we devise **selective L₂ loss**.
- ❖ **Identify** pixels where **predicted labels may be wrong** but the world model **decoder is correct**, **ignoring L₂ loss** for such pixels to avoid providing wrong signals.

$$\text{mask}_{\text{estimate}} = \text{mask}_{\text{SAM}} \vee \neg \text{mask}_{\text{SD}}$$



4. Experimental Results



- ❖ **Dreamer*** is trained in a **clean** environment whereas the other methods are trained in a **distracting** environment.
- ❖ Overall, **SD_{GT}** matches **Dreamer***, and **SD_{approx.}** eventually reaches **SD_{GT}** while **Dreamer** falls short.

- ❖ **Reconstruction-free methods** take **lots of samples** to converge.
- ❖ In **sparse reward** scenario, **only our method** achieves compelling performance.

5. References

[1] Hafner et al. Mastering diverse domains through world models. arXiv preprint, 2023.

[2] Hansen et al. Td-mpc2: Scalable, robust world models for continuous control. ICLR, 2024.

[3] Deng et al. Dreamerpro: Reconstruction-free model-based reinforcement learning with prototypical representations. ICML, 2022.

[4] Zhang et al. Personalize segment anything model with one shot. arXiv preprint, 2023.