Learning Task Informed Abstractions

Xiang Fu*, Ge Yang*, Pulkit Agrawal, Tommi Jaakkola

xiangfu.co/tia

The success of model-based agents in visual domains

DeepMind Control

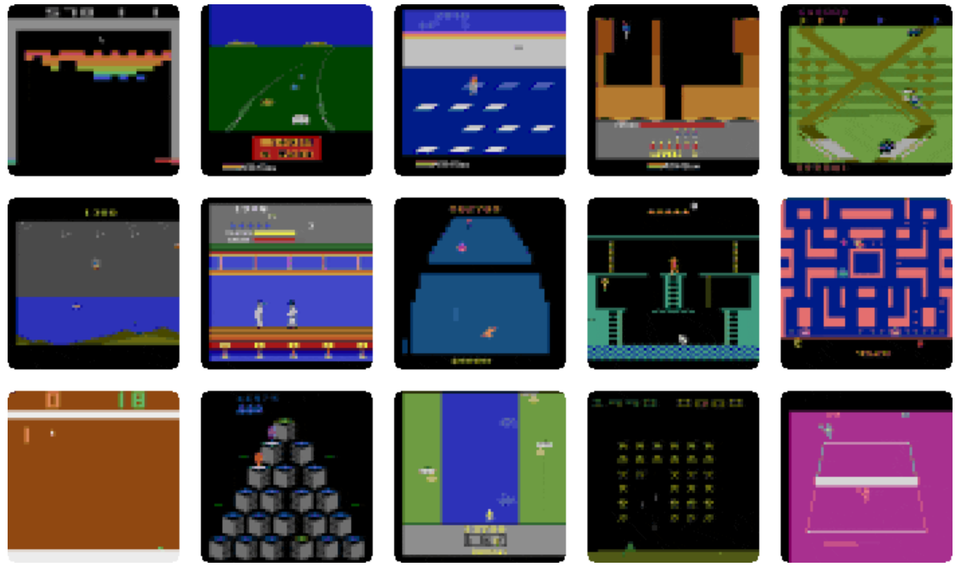

ATARI

Image source: https://dreamrl.github.io/, https://danijar.com/project/dreamerv2/

The success of model-based agents in visual domains

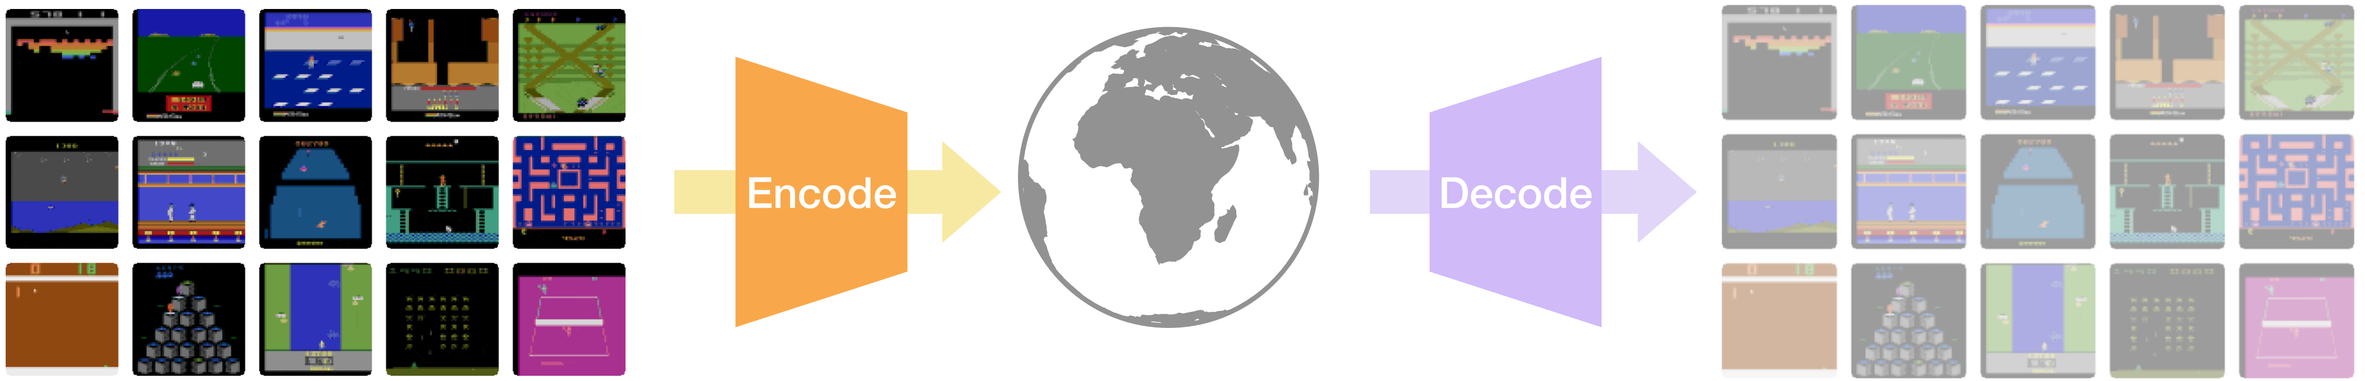

Observation

World Model

Reconstruction

https://danijar.com/project/dreamerv2/

Model-based agents suffer from complex visual distractors

“Cheetah Run”

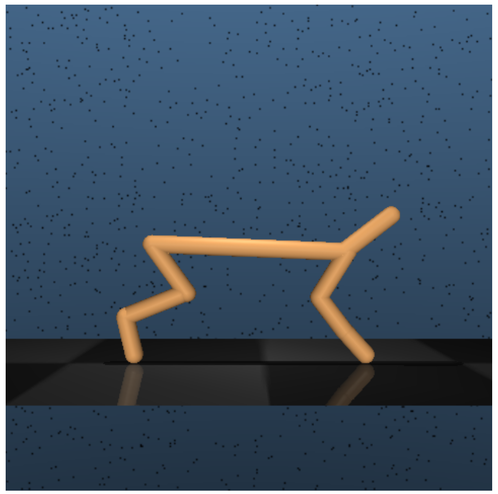

“With Distractions”

Model-based agents suffer from complex visual distractors

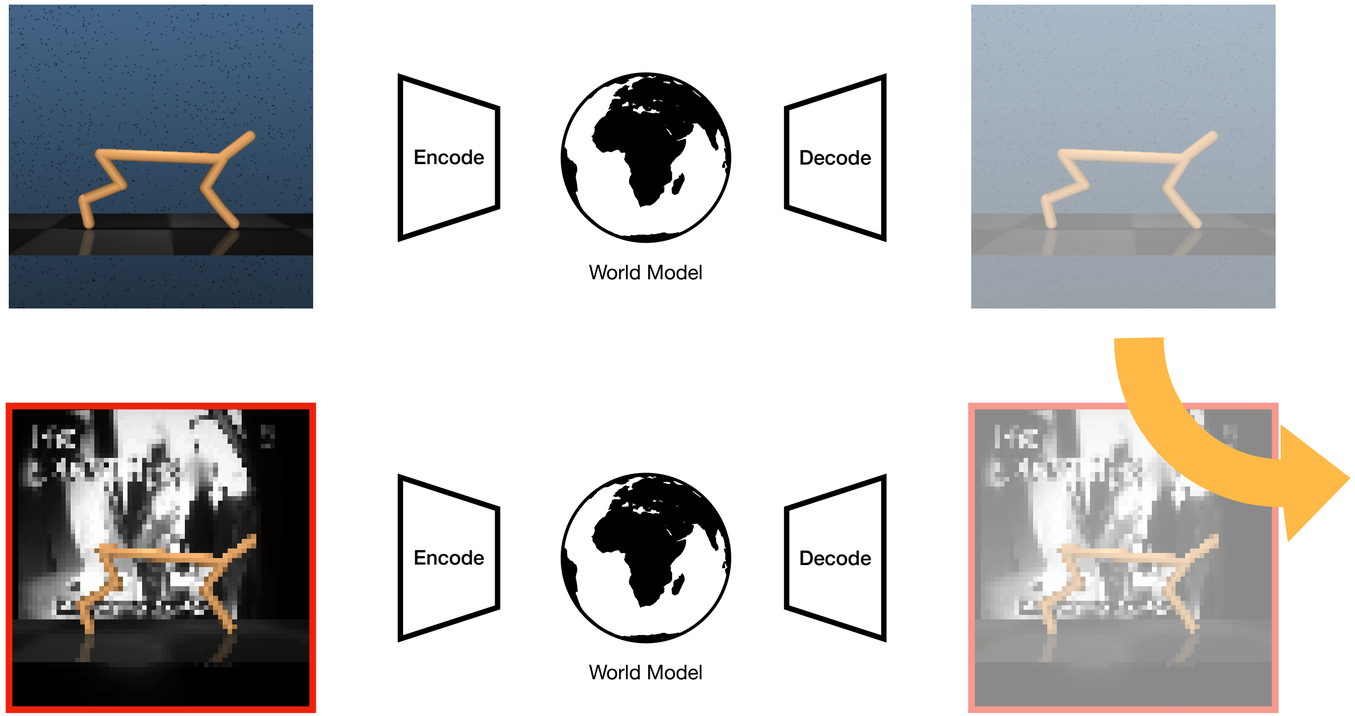

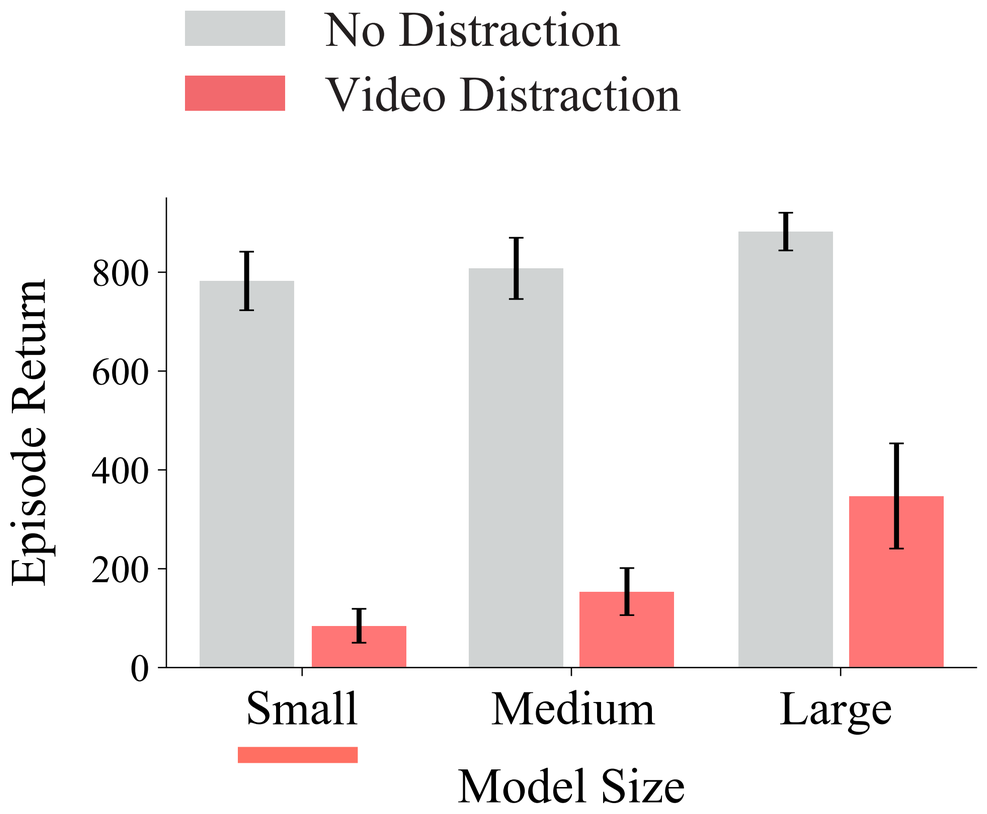

Model capacity is consumed by the non-informative visual distractions

Part I: model learning

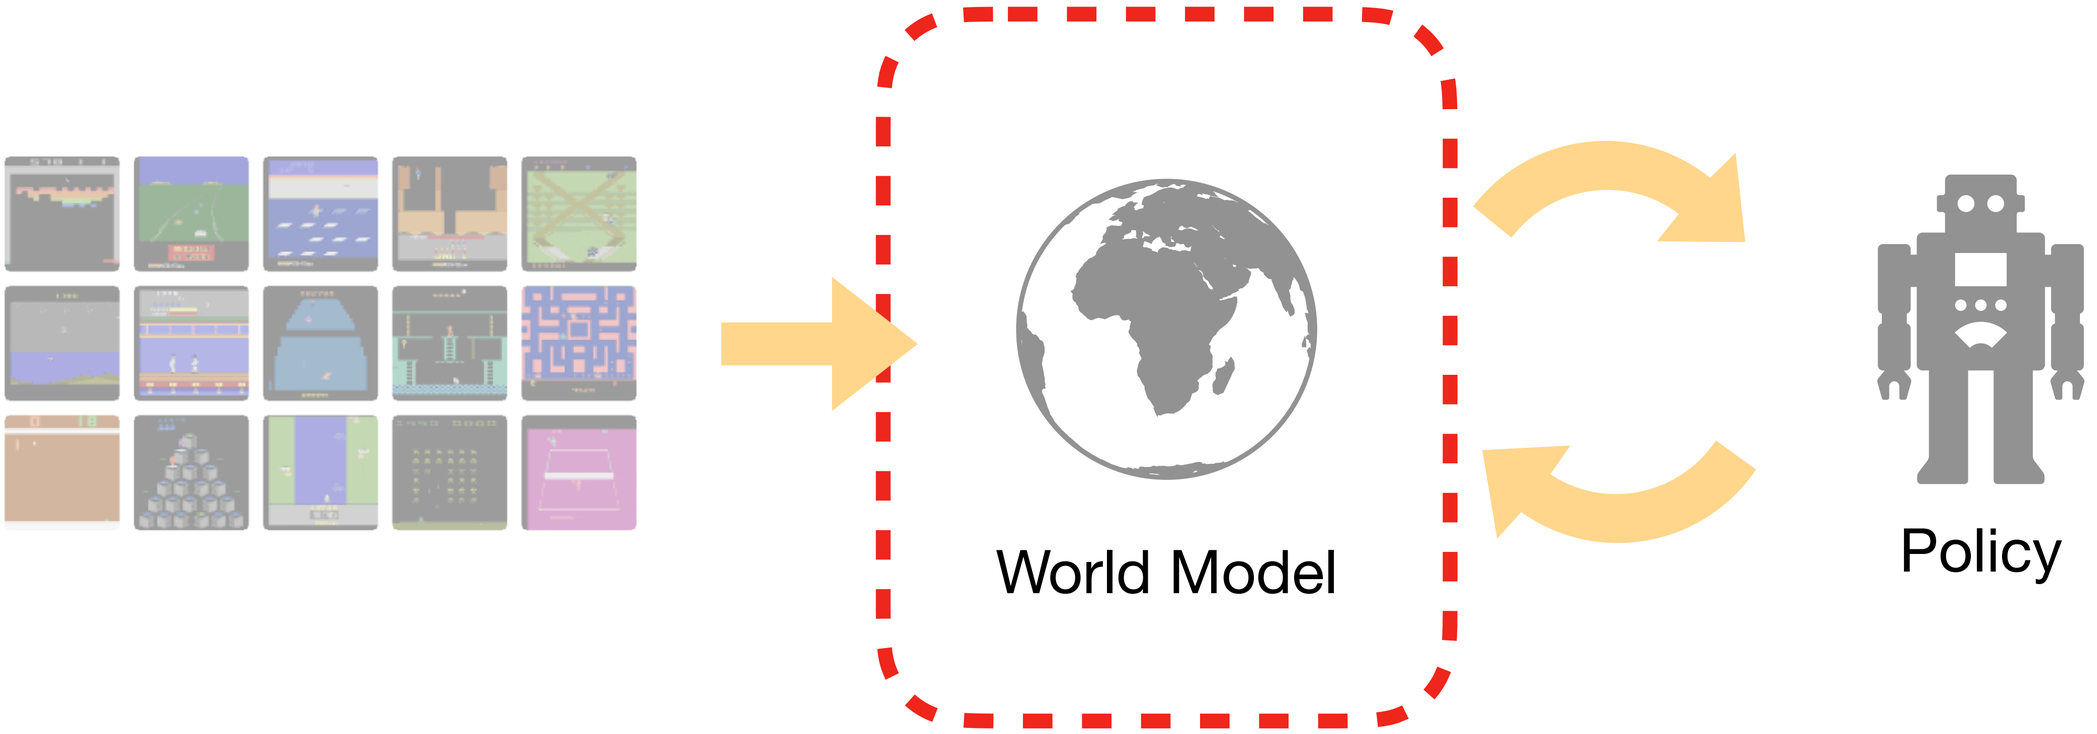

https://danijar.com/project/dreamerv2/

Part I: model learning

Incentivize s_t^+ to predict rewards

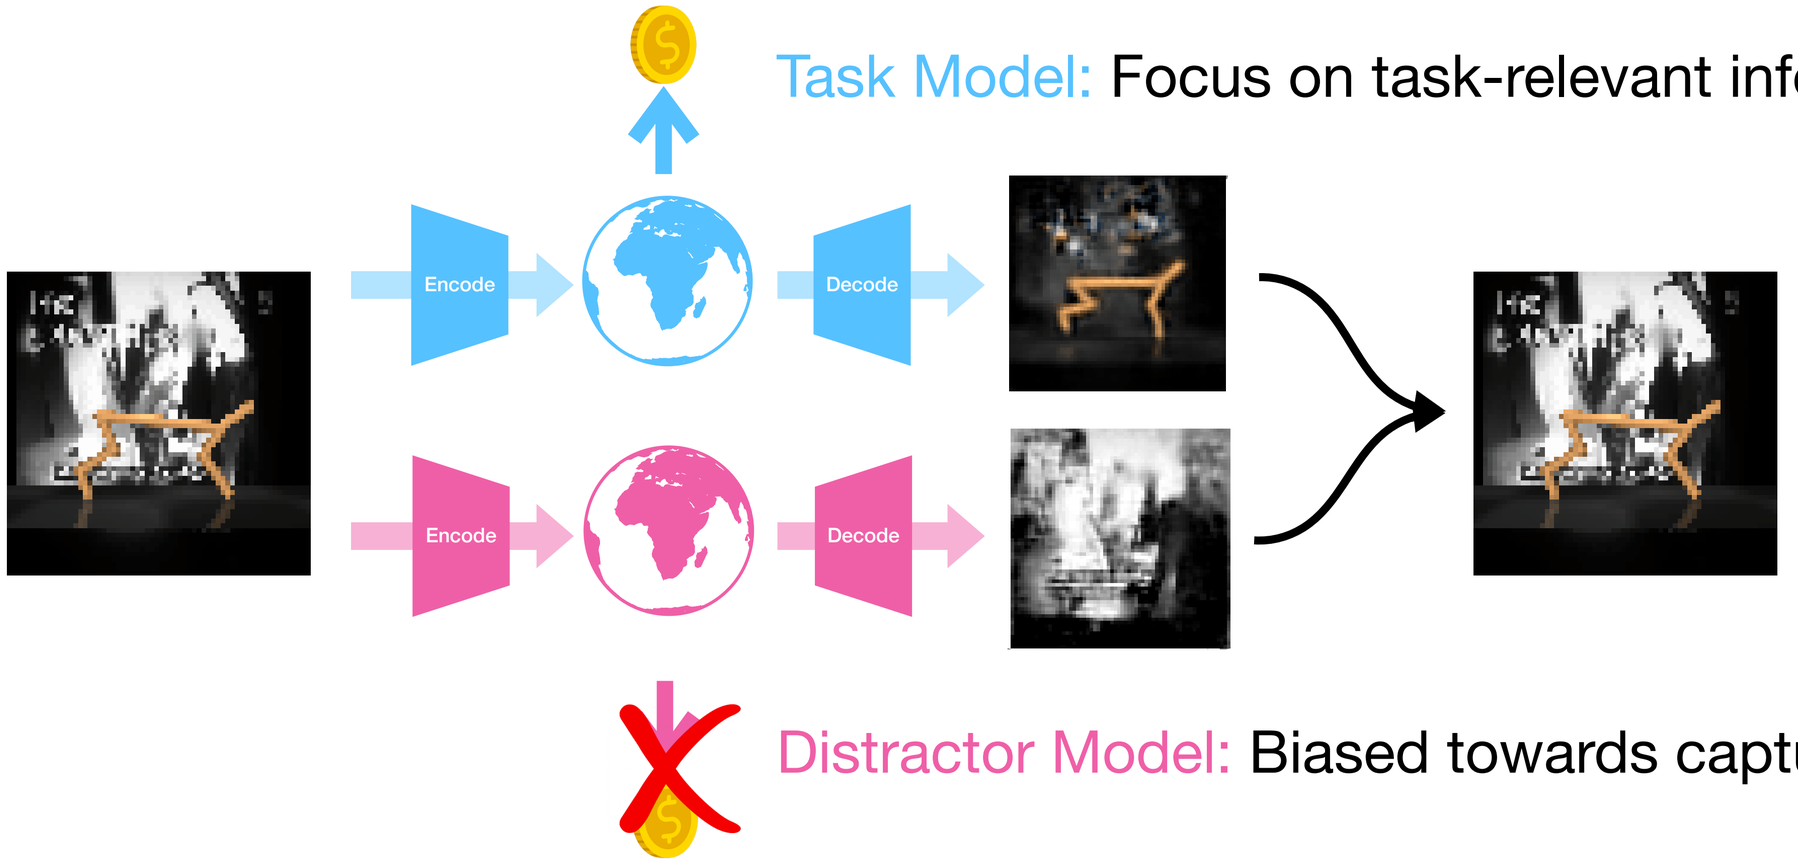

Task Model: Focus on task-relevant information

Joint Reconstruction:
$\max\ p(o_t \mid s_t^+, s_t^-)$

Distractor Model: Biased towards capture distractors

Constrain s_t^- NOT to predict rewards

Part II: policy learning

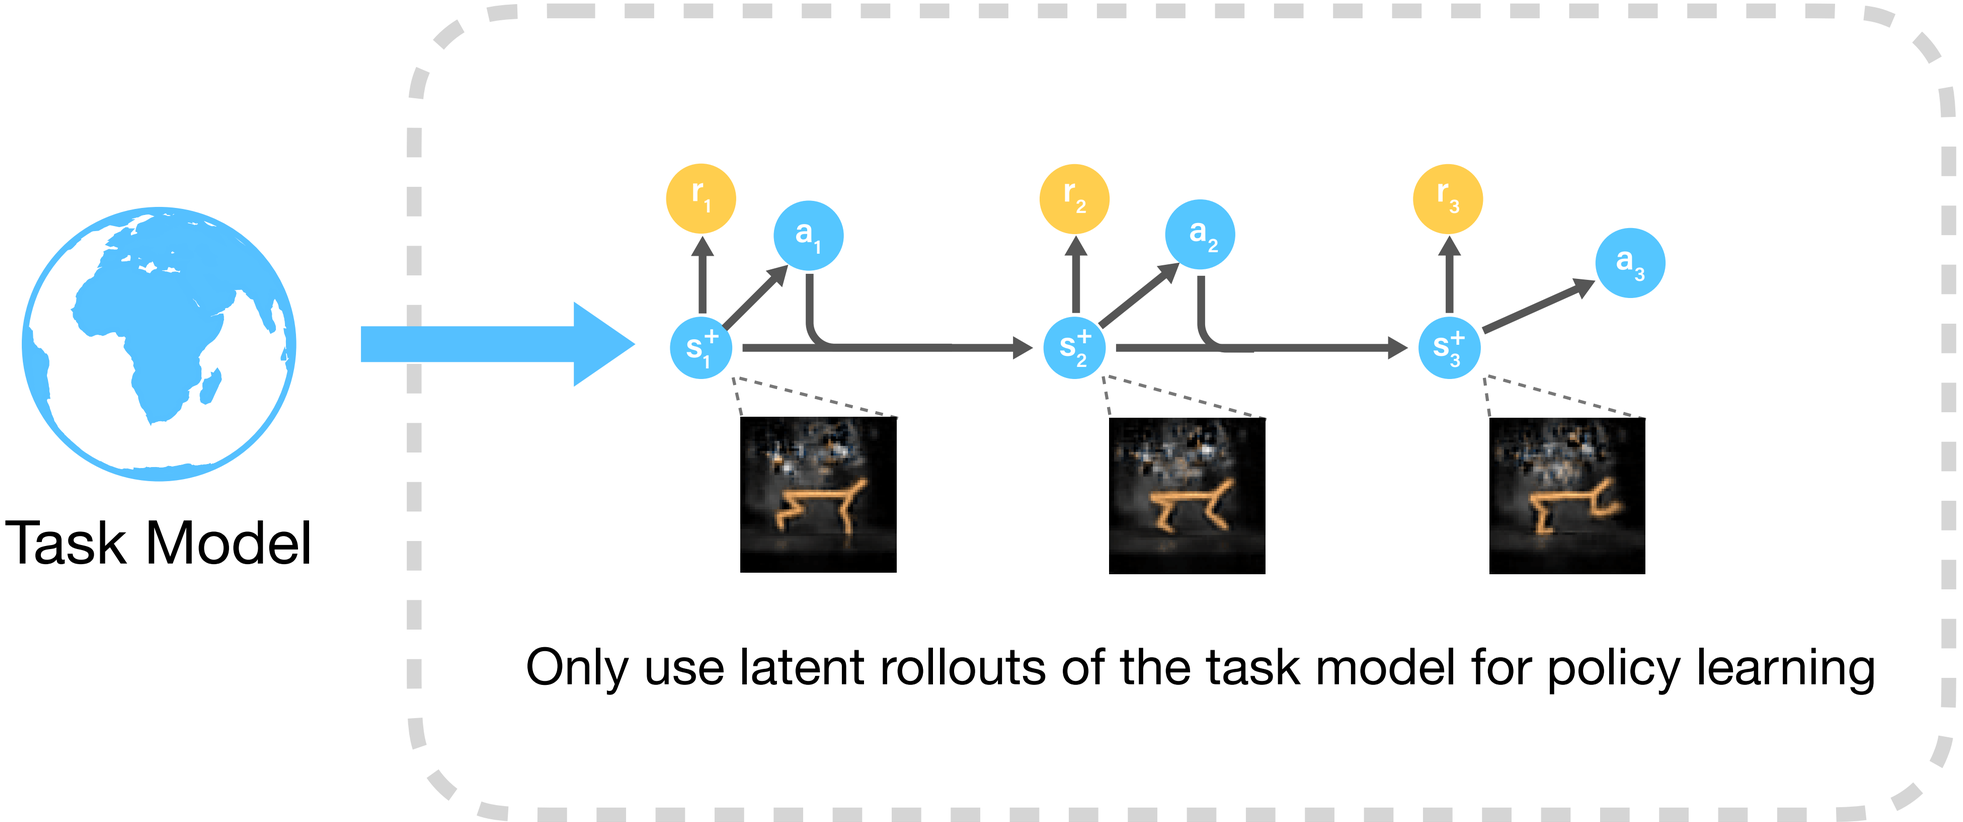

Learned representation of task-informed abstractions

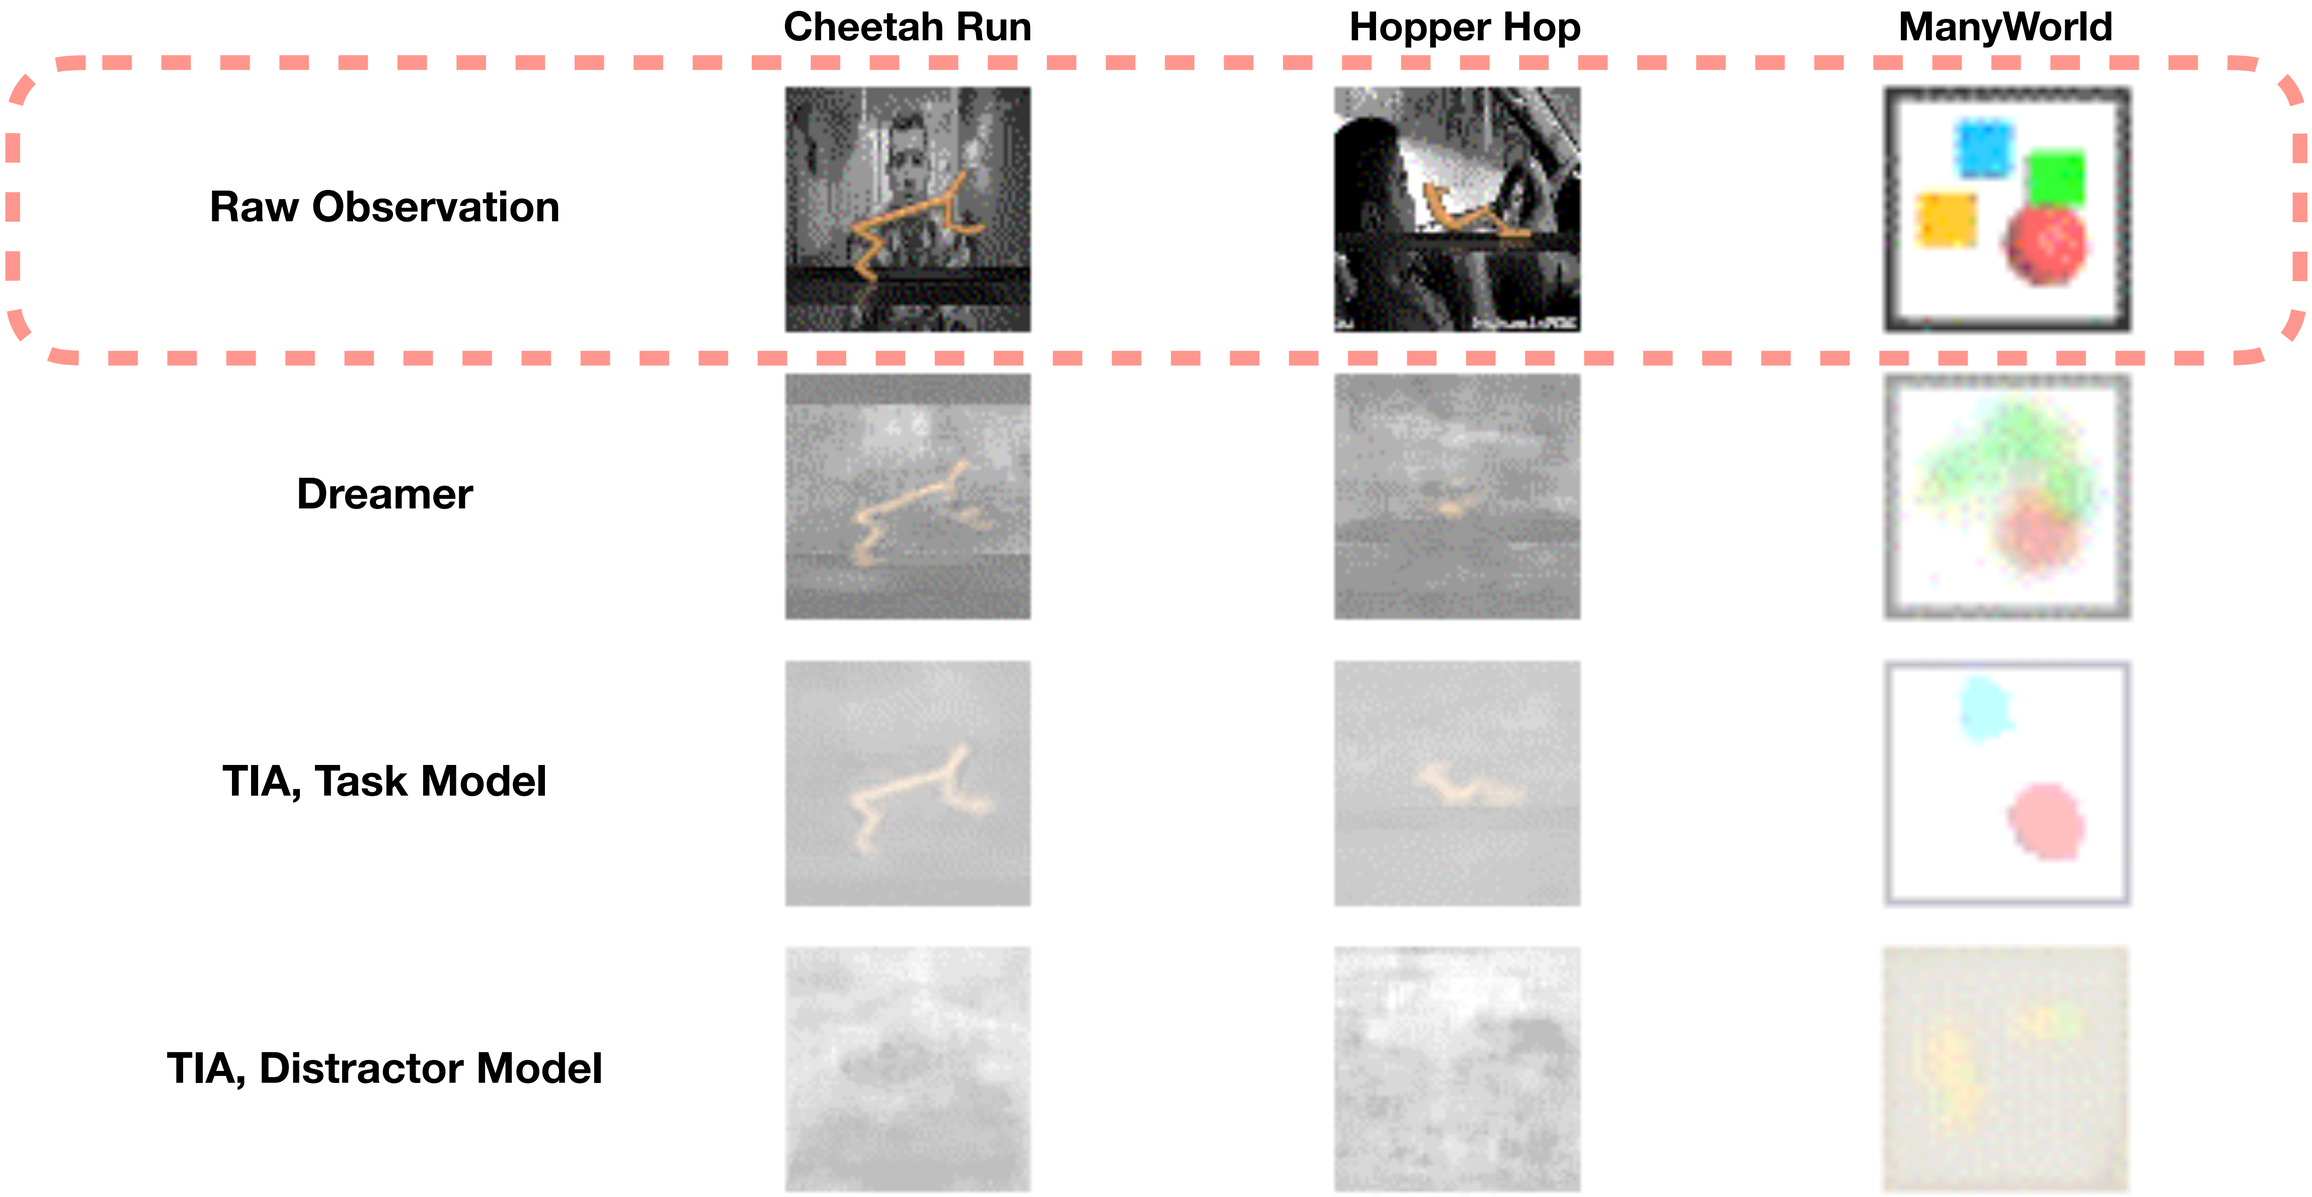

Performance in complex visual domains: DMC

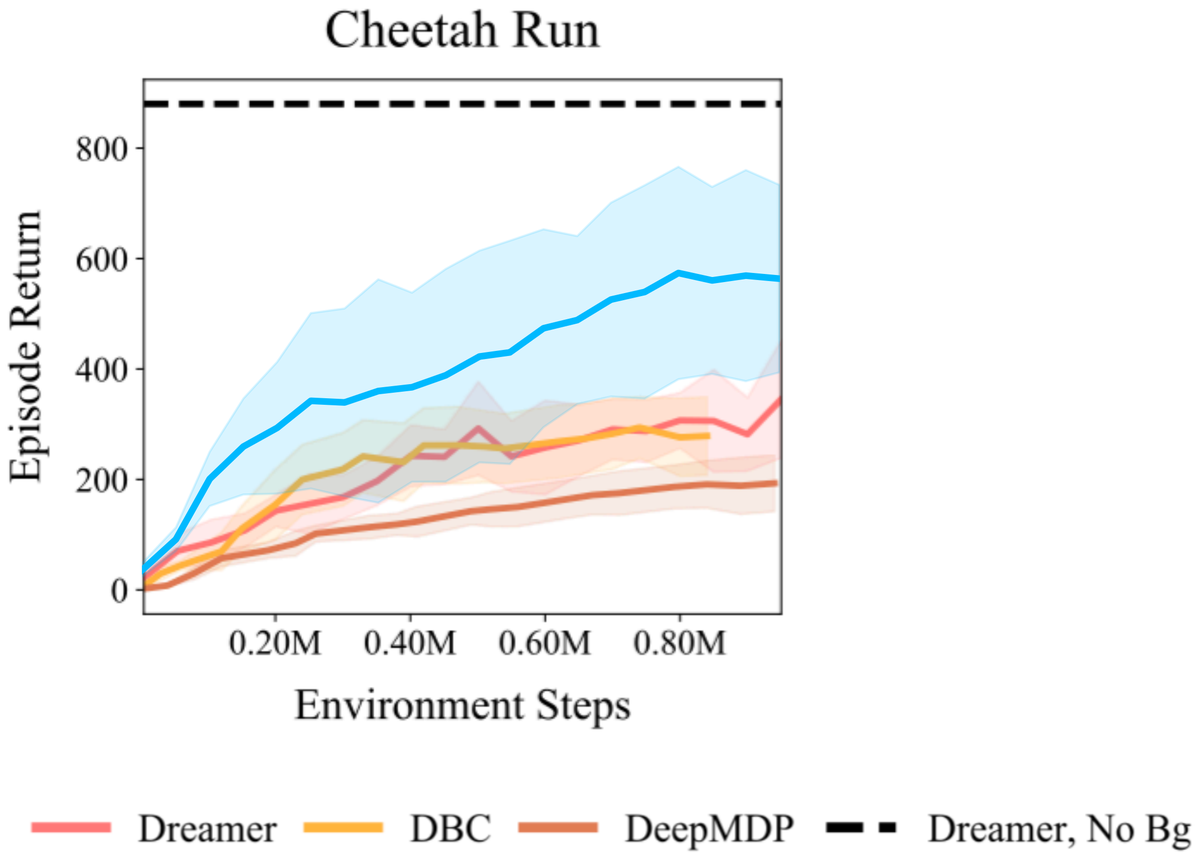

Performance in complex visual domains: DMC

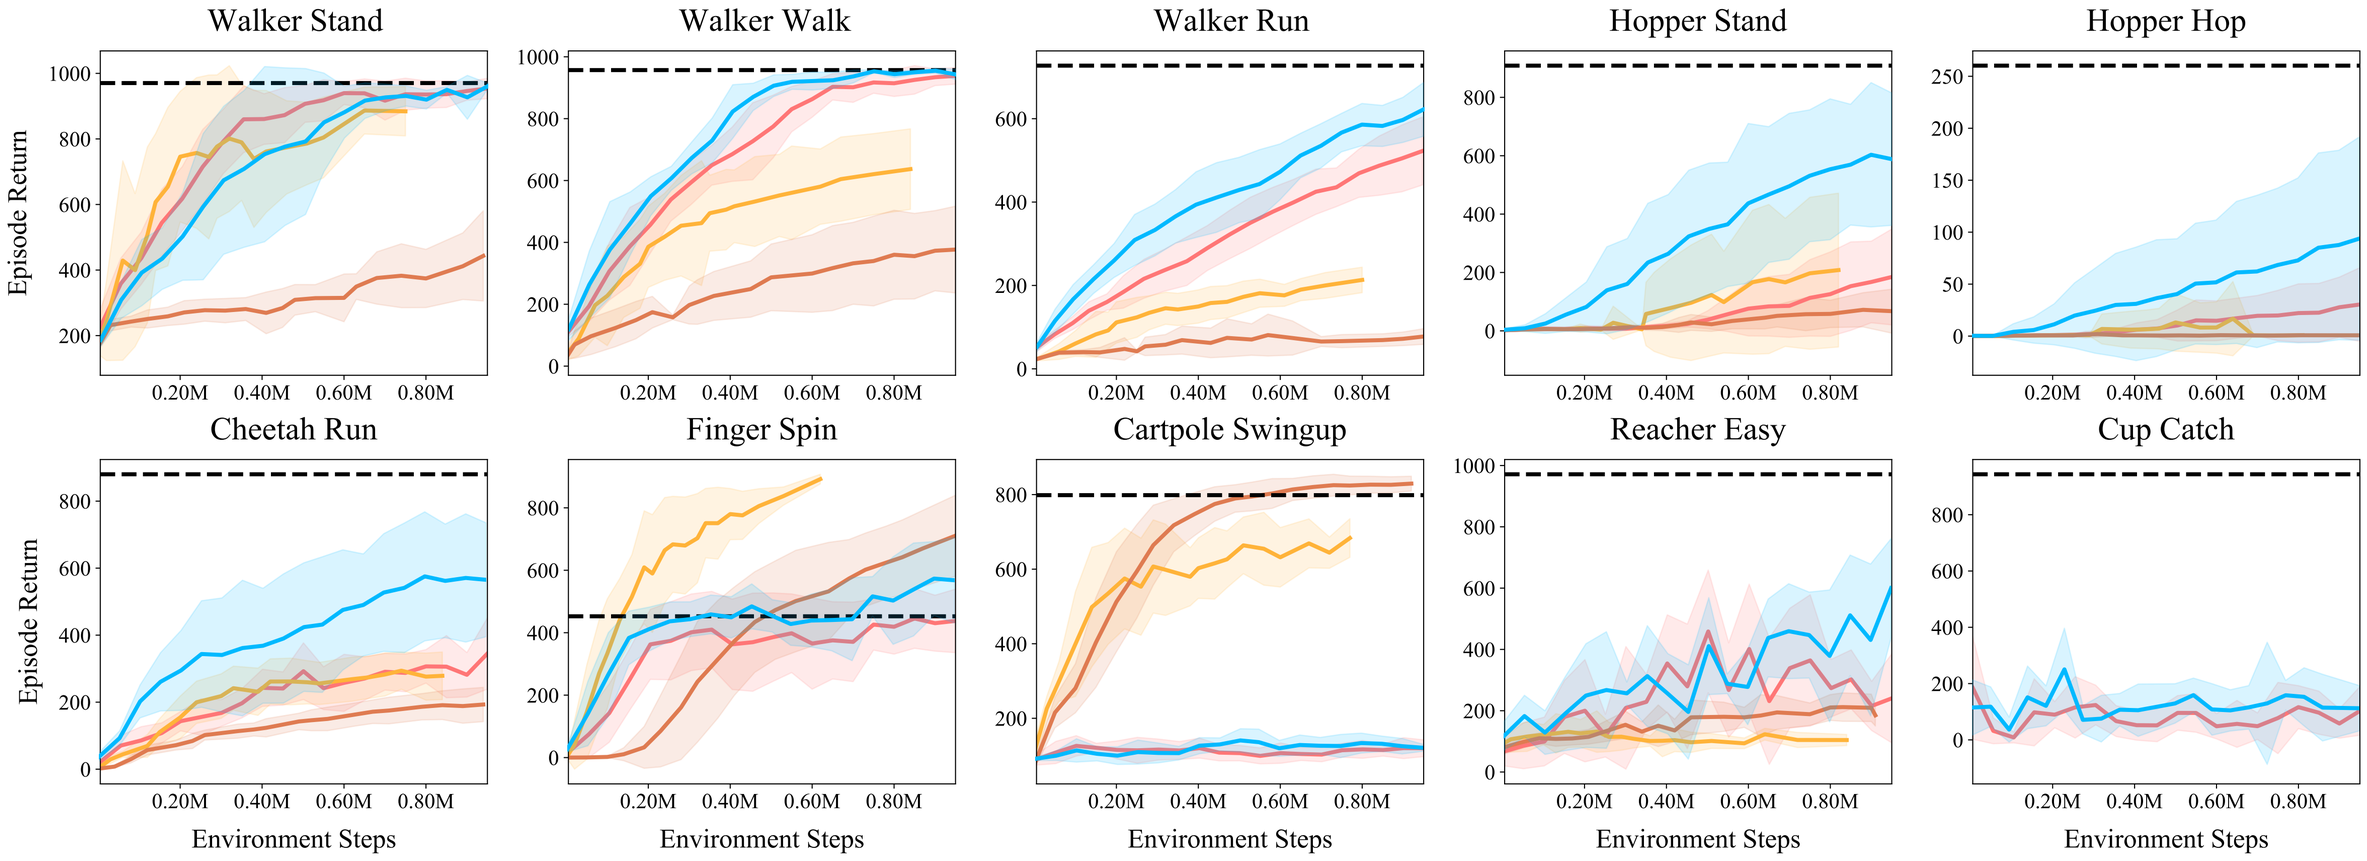

Performance in complex visual domains: Atari

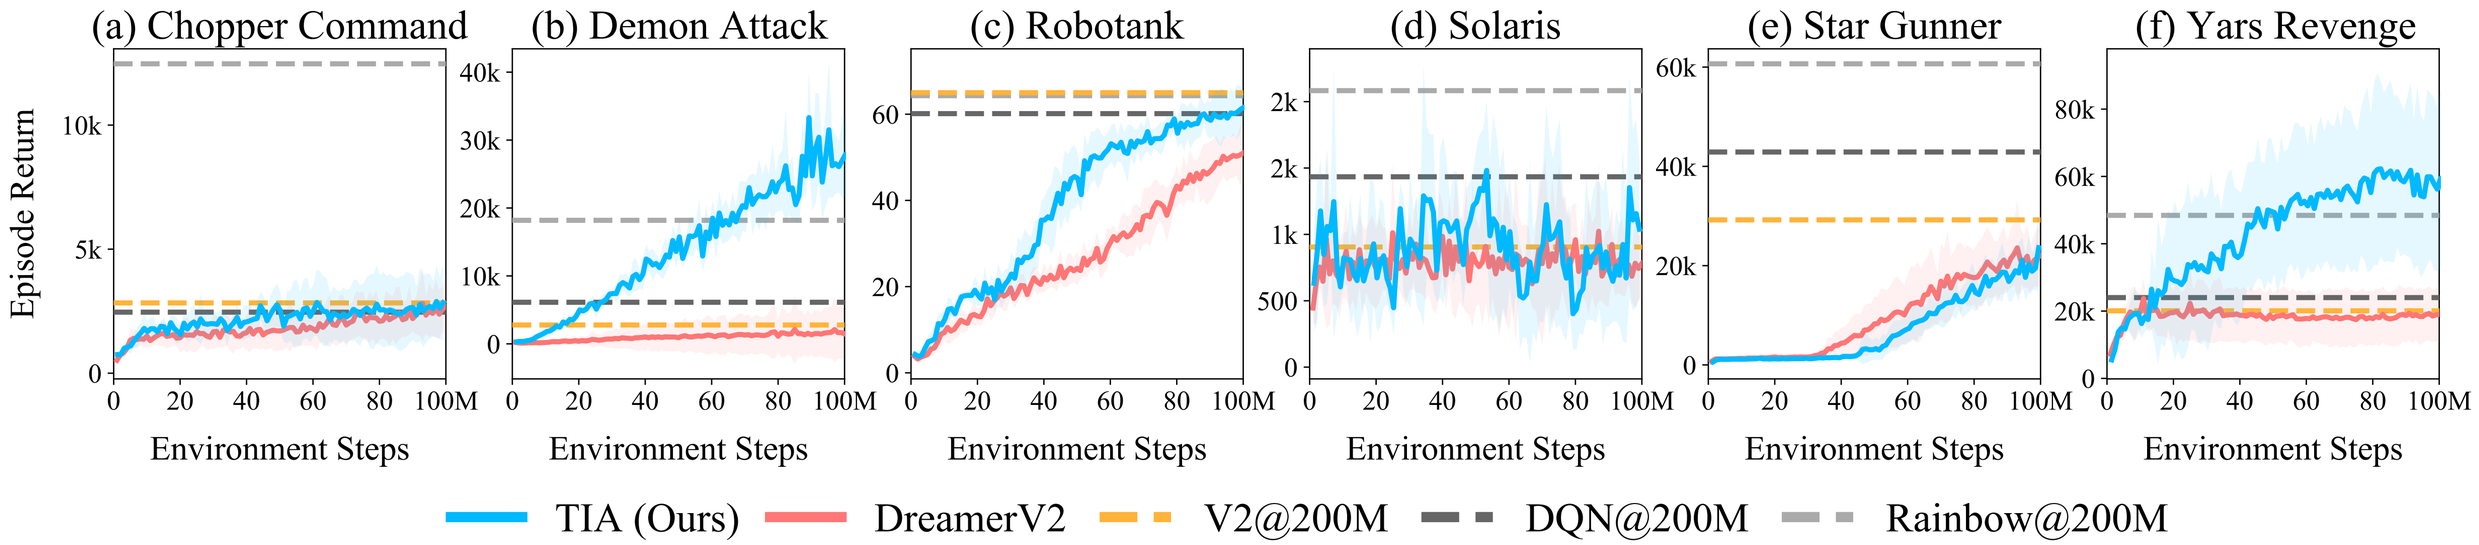

Performance in complex visual domains: Atari

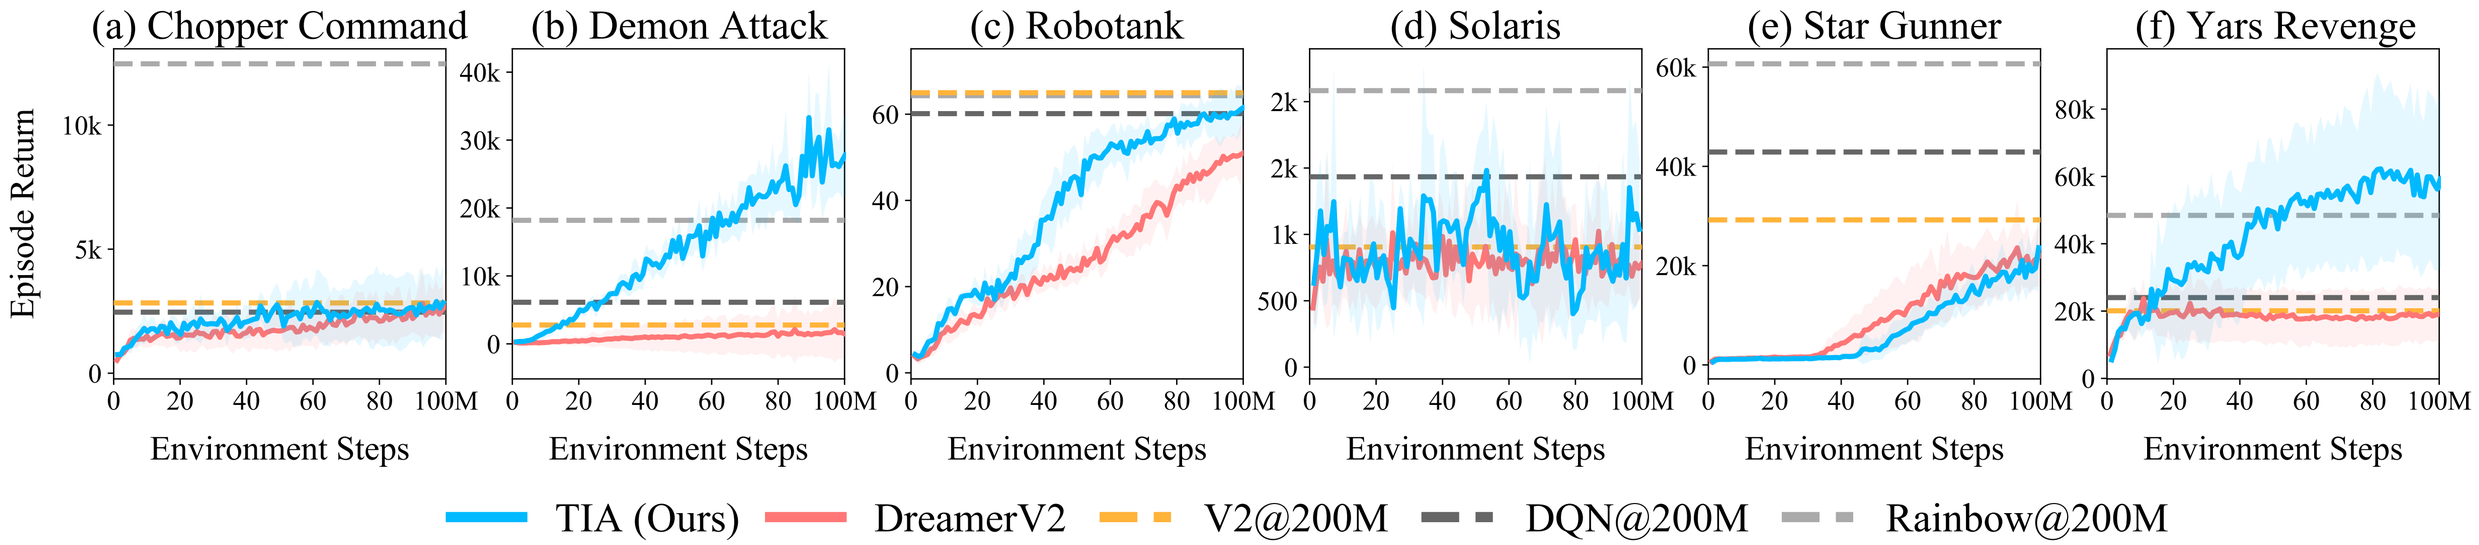

Learning Task Informed Abstractions

Xiang Fu*, Ge Yang*, Pulkit Agrawal, Tommi Jaakkola

Paper & code available at: xiangfu.co/tia