

1 We introduce Dreamer

- We present Dreamer, an RL agent that solves visual control tasks purely by latent imagination inside of a world model.
- $\boxed{2}$

To efficiently learn behaviors, we backpropagate multi-step value gradients through imagined latent transitions.

3

Dreamer learns off-policy and is applicable to both continuous and discrete actions and early episode ends.

Dreamer exceeds top model-free agents in final performance, data-efficiency, and wall-clock time.

4 Learning Behaviors in Imagination

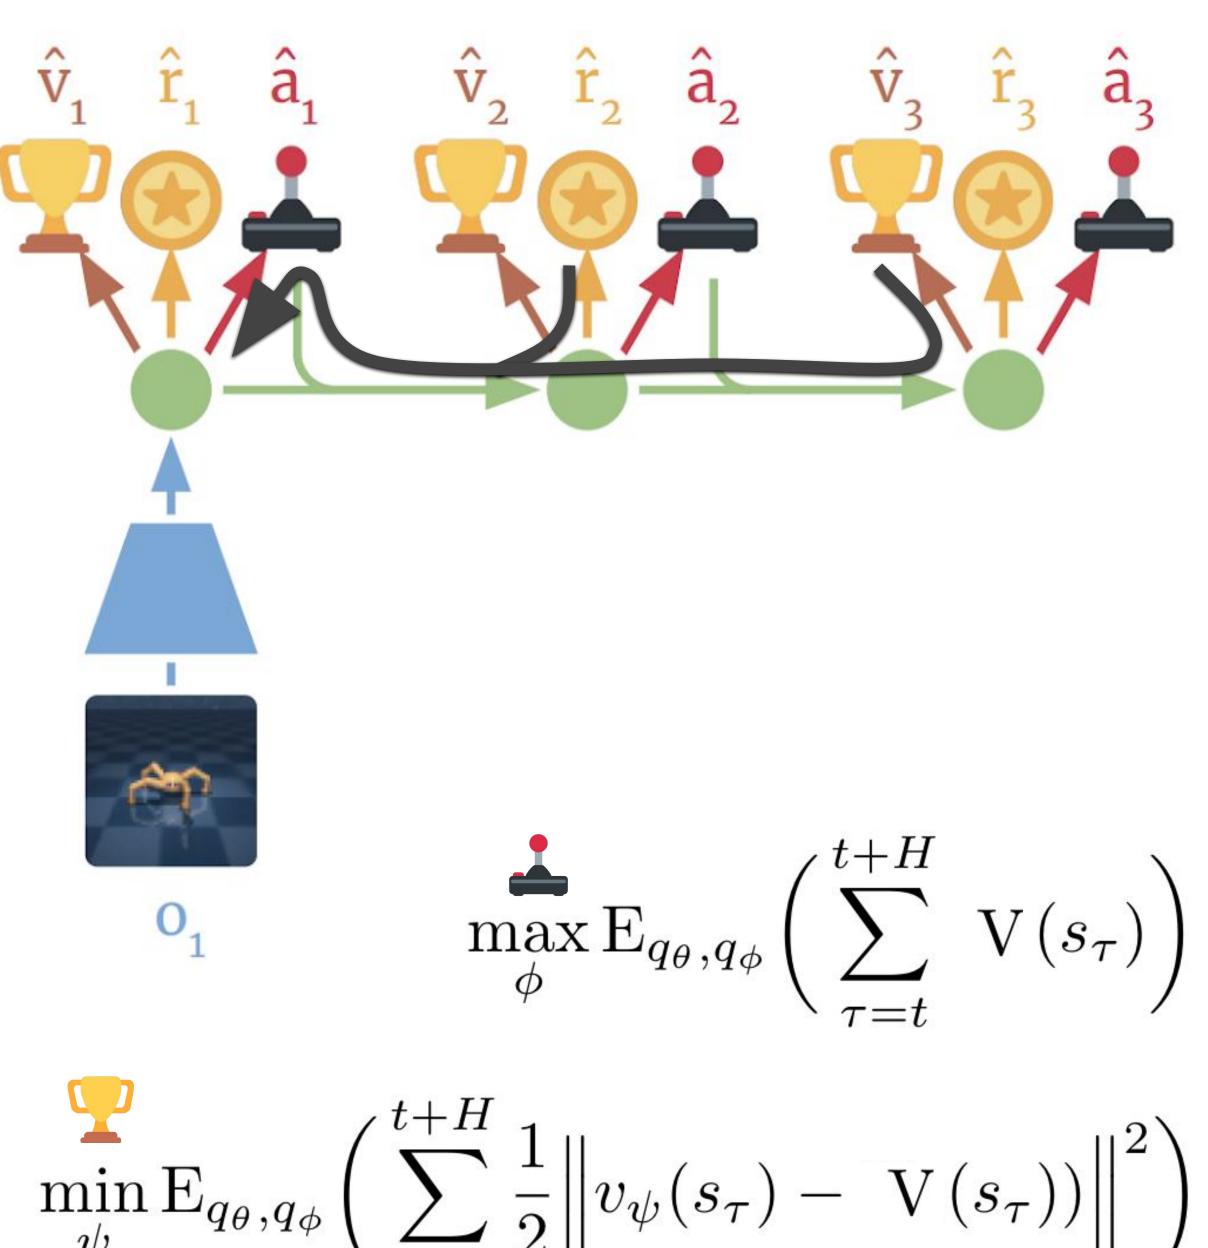
Dreamer learns to predict actions and values from the compact model states.

Start imagined trajectory from each latent state of the current data batch.

Value model optimizes Bellman consistency for current action model.

Action model maximizes values by backpropagating through latent transitions.

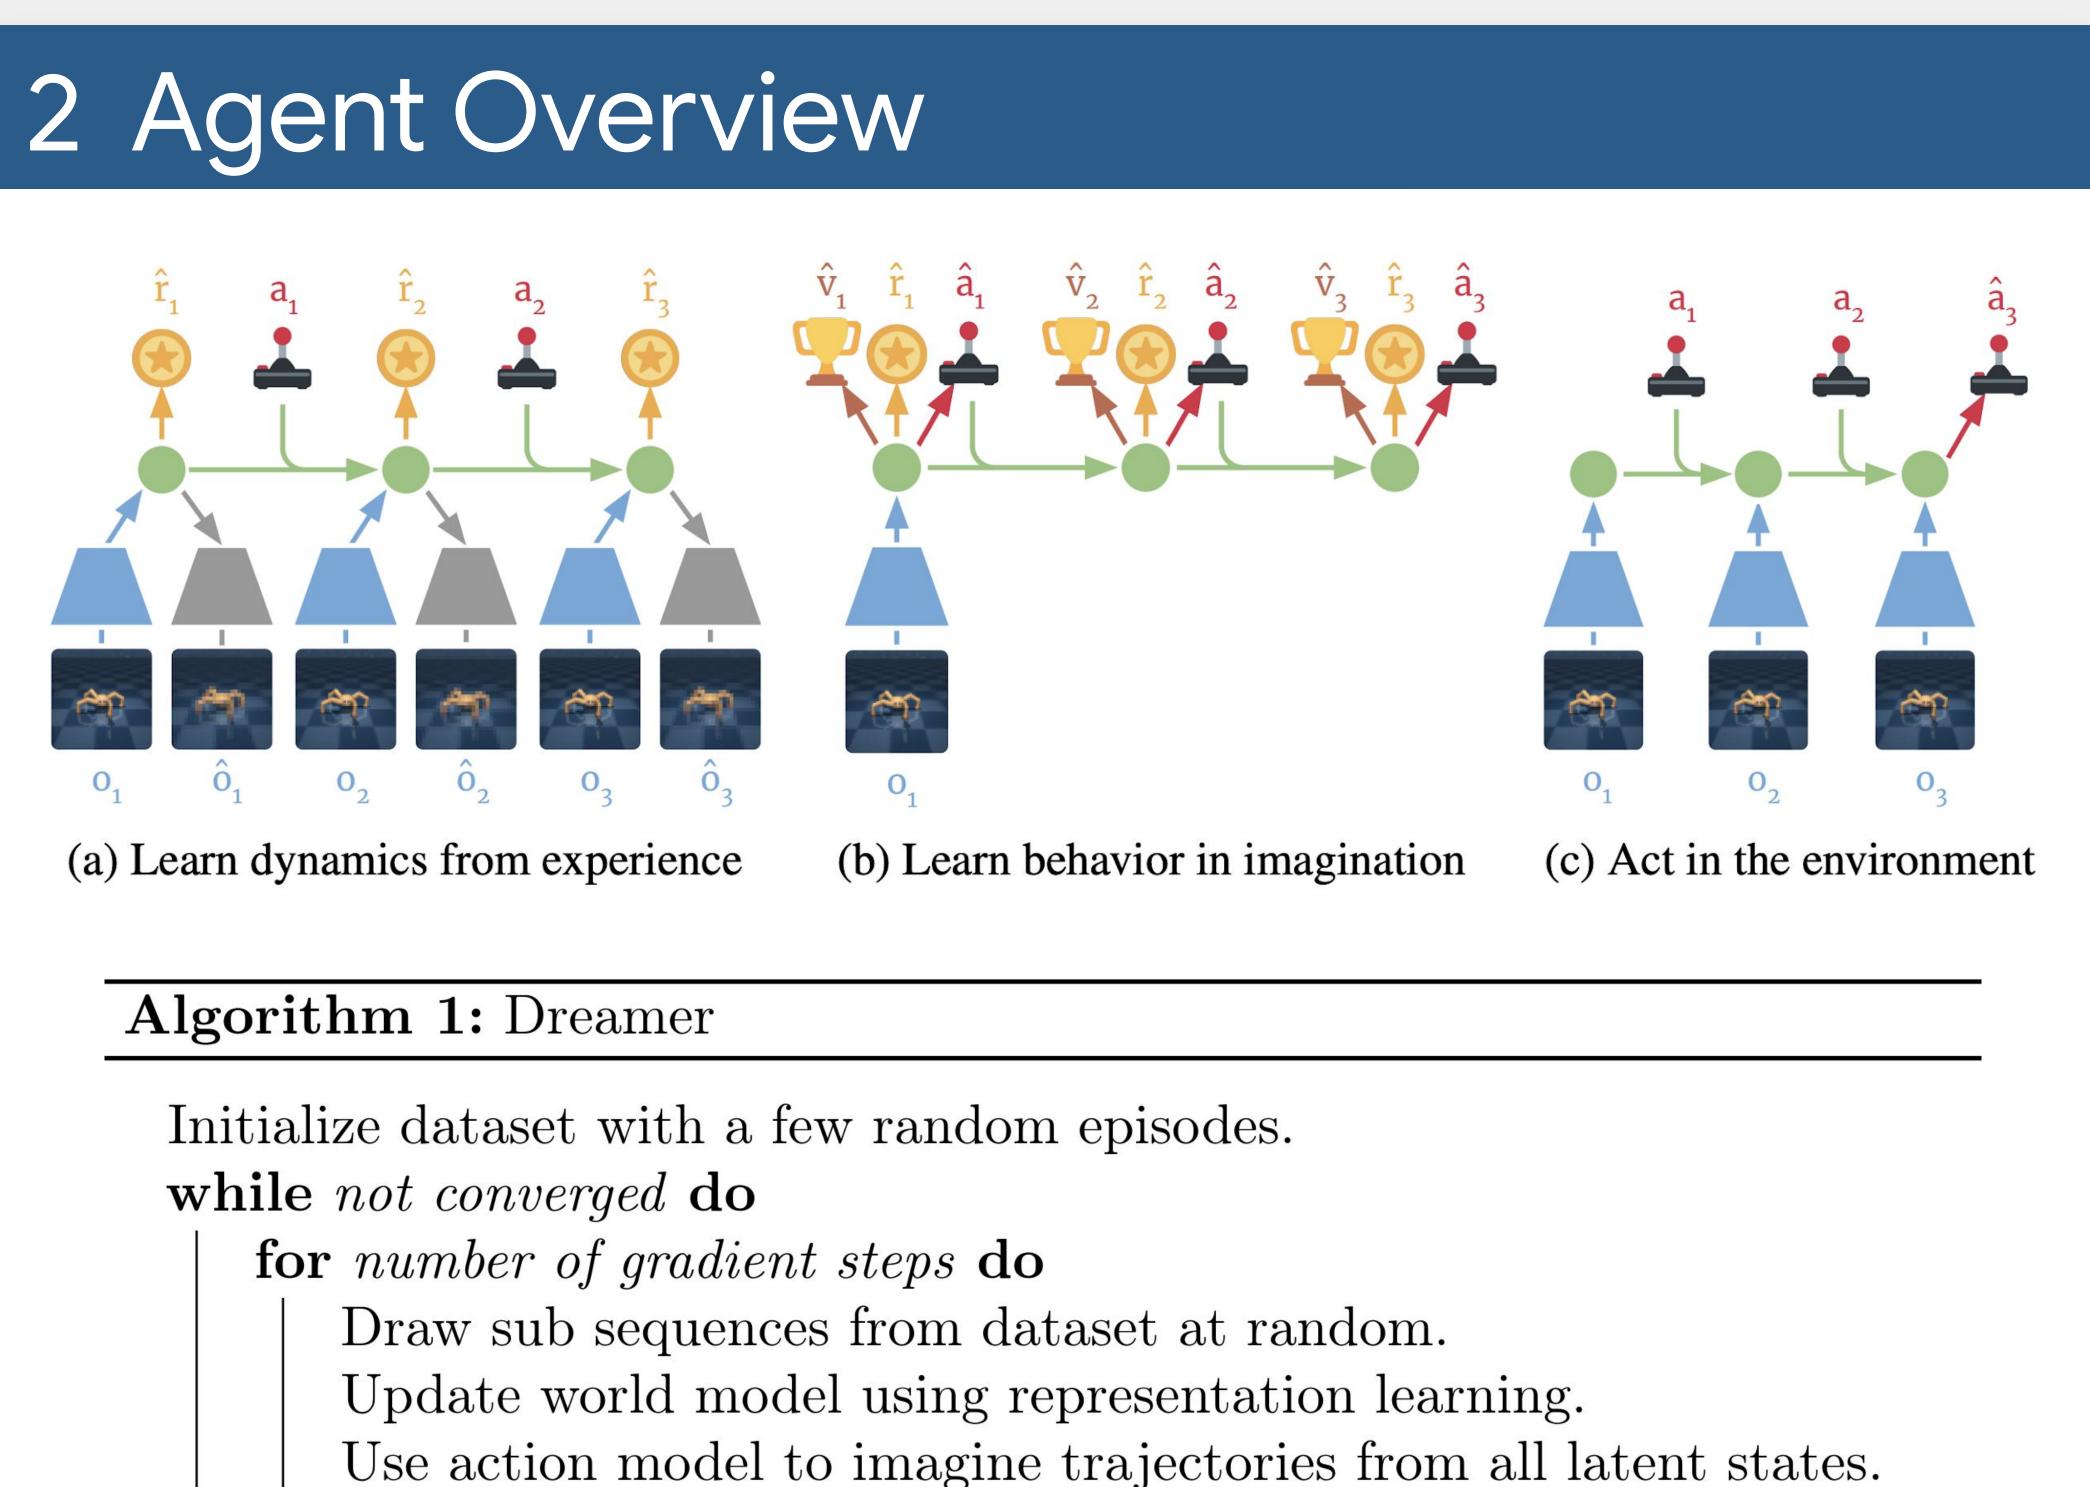
Reparam. gradients for latent states and actions, ST for discrete actions.



 $\mathcal{V}(s_{\tau}) \stackrel{\text{e.g.}}{=} r(s_{\tau}) + \gamma r(s_{\tau+1}) + \gamma^2 v_{\psi}(s_{\tau+2})$

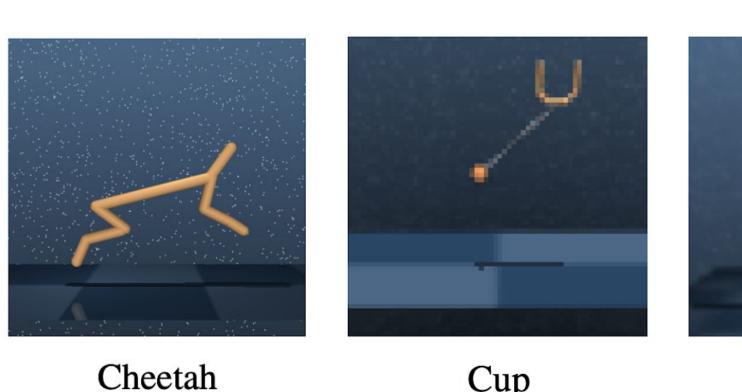
Dream to Control Learning Behaviors by Latent Imagination Danijar Hafner¹², Timothy Lillicrap³, Jimmy Ba², Mohammad Norouzi¹ ¹Google Brain ²University of Toronto ³DeepMind





 $V(s_{\tau})) \Big\|^2 \Big)$ $v_{\psi}(s_{\tau})$

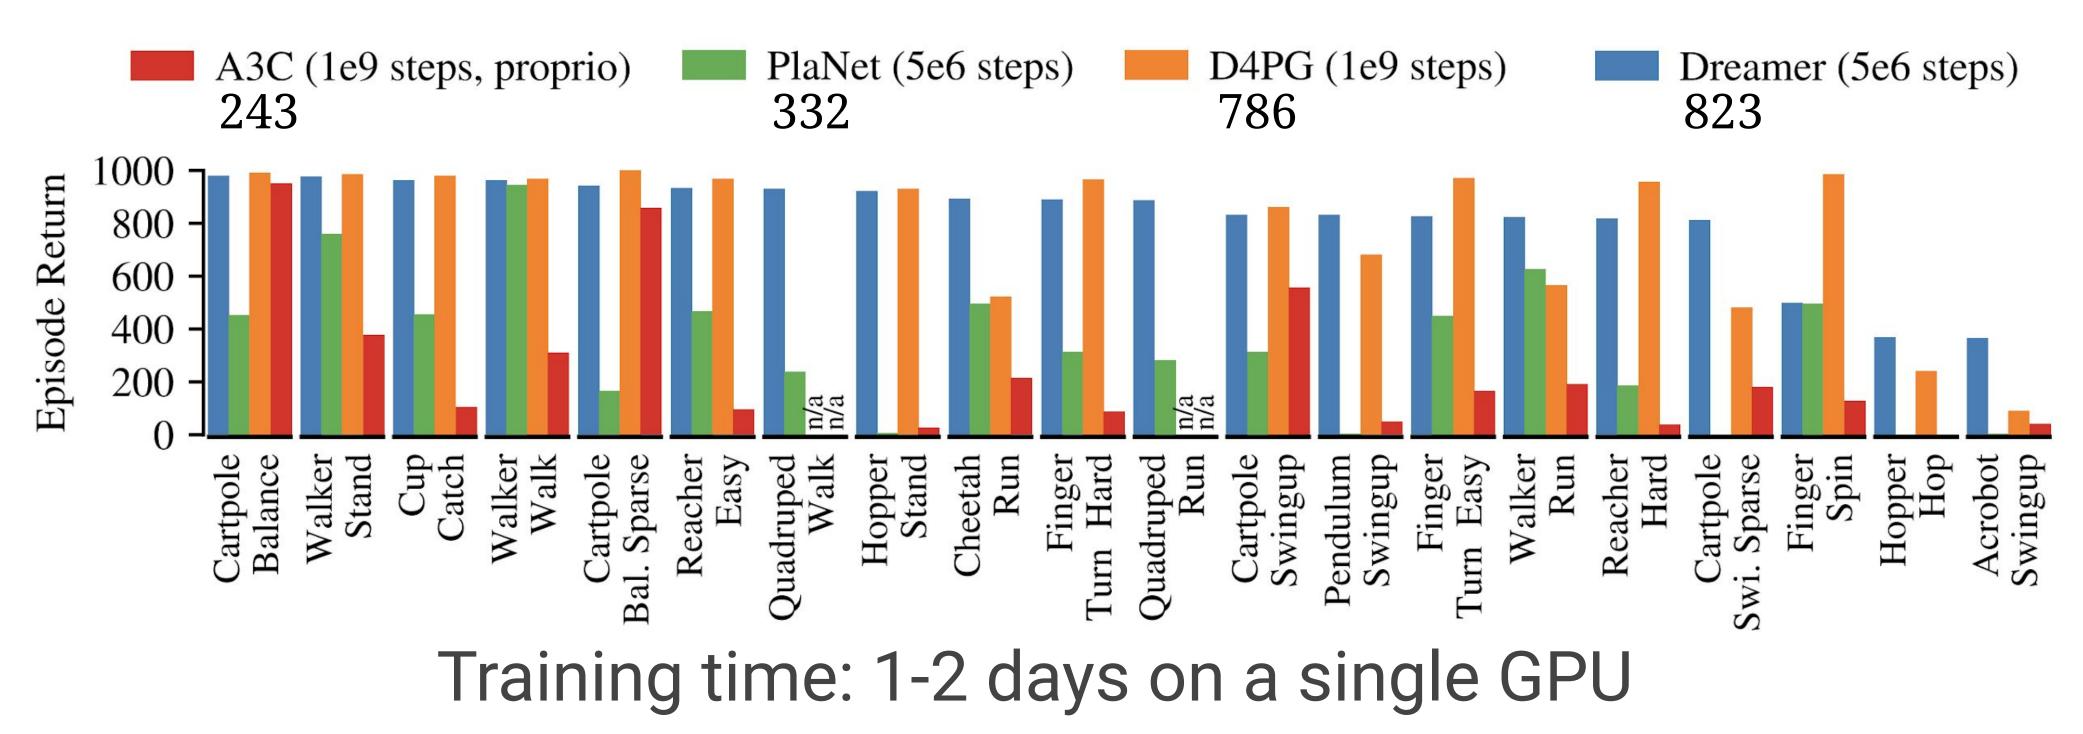
5 Comparison to Existing Methods



Acrobot

20 visual control tasks posing challenges incl. contact dynamics, sparse rewards, many degrees of freedom, and 3D environments.

Dreamer outperforms the previous best D4PG agent in final performance, sample-efficiency, and wall-clock time. We use the same hyper parameters across all tasks.



- Update value model for Bellman consistency on imagined rewards.
- Update action model to maximize imagined values.
- Collect episode to grow the dataset using the action model.







Walker

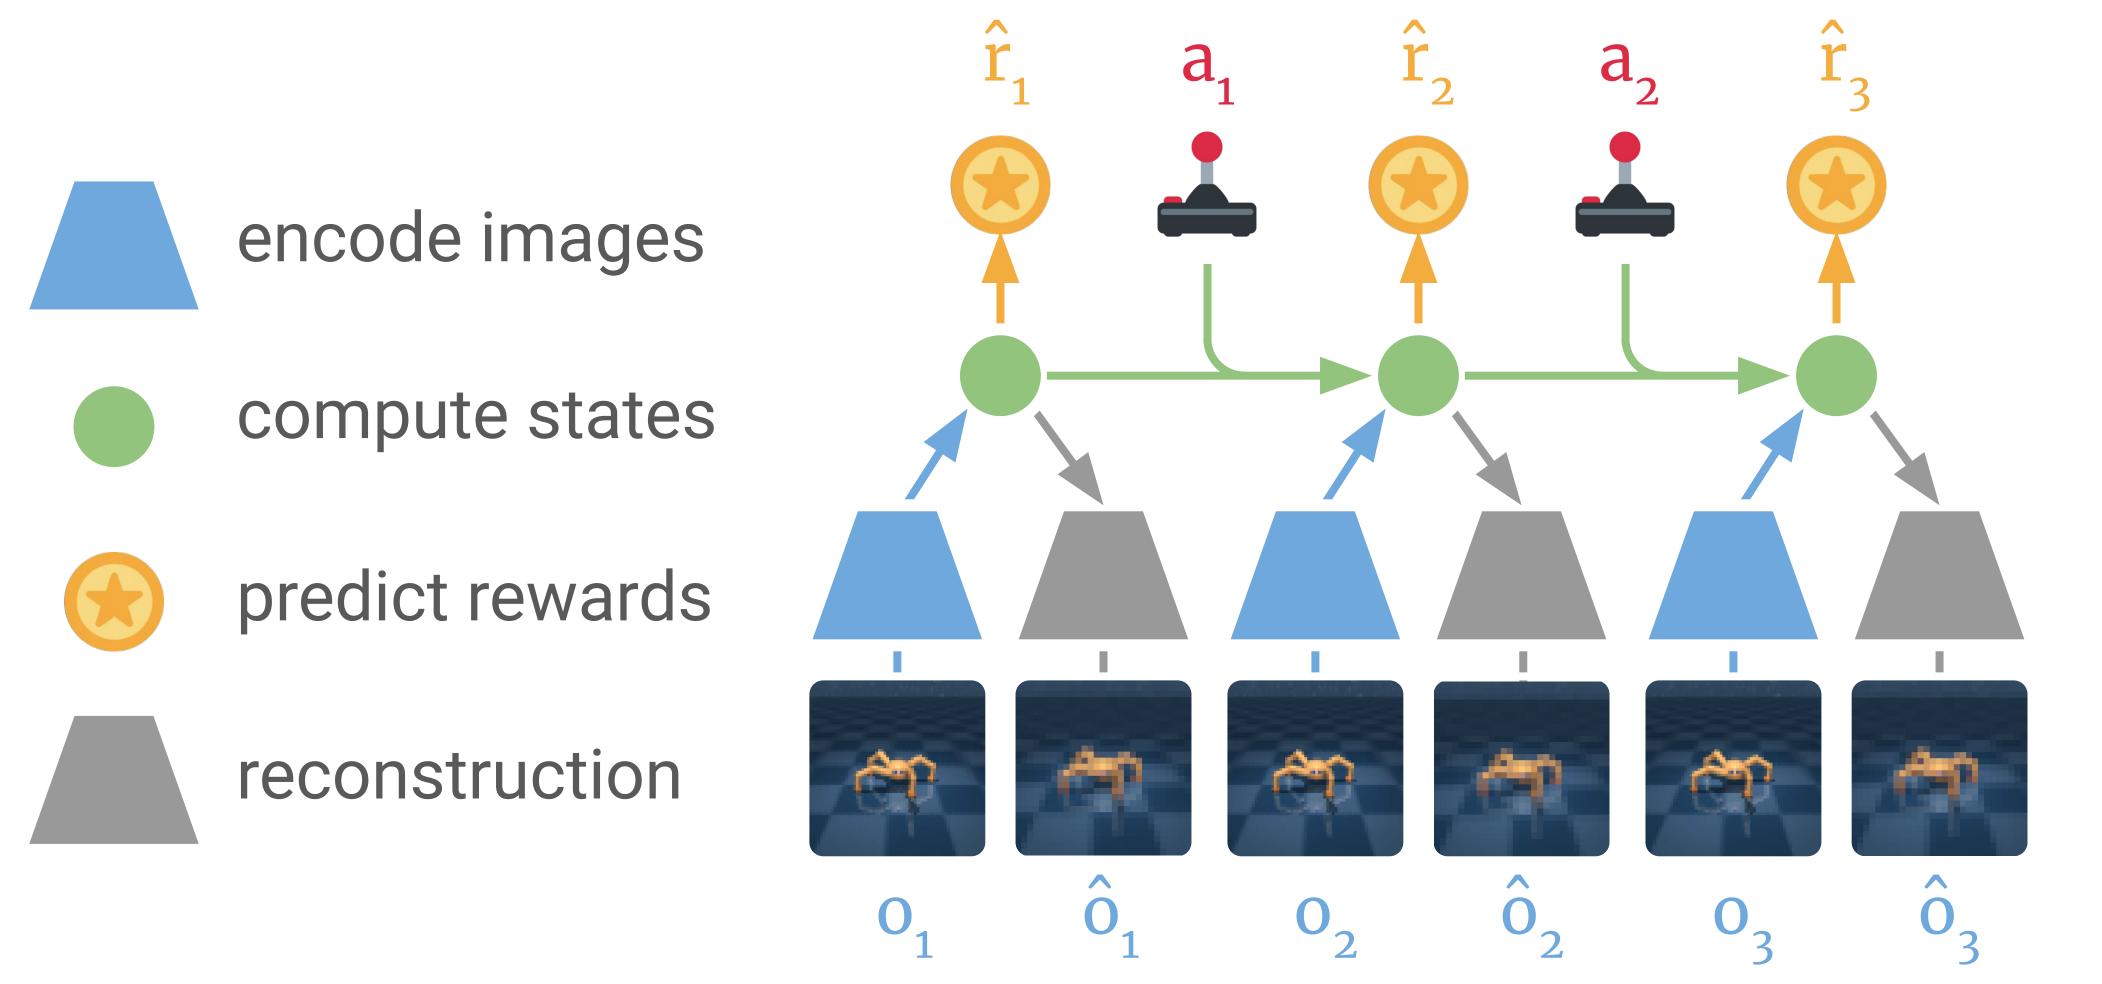
Quadruped

3 Learning Latent Dynamics

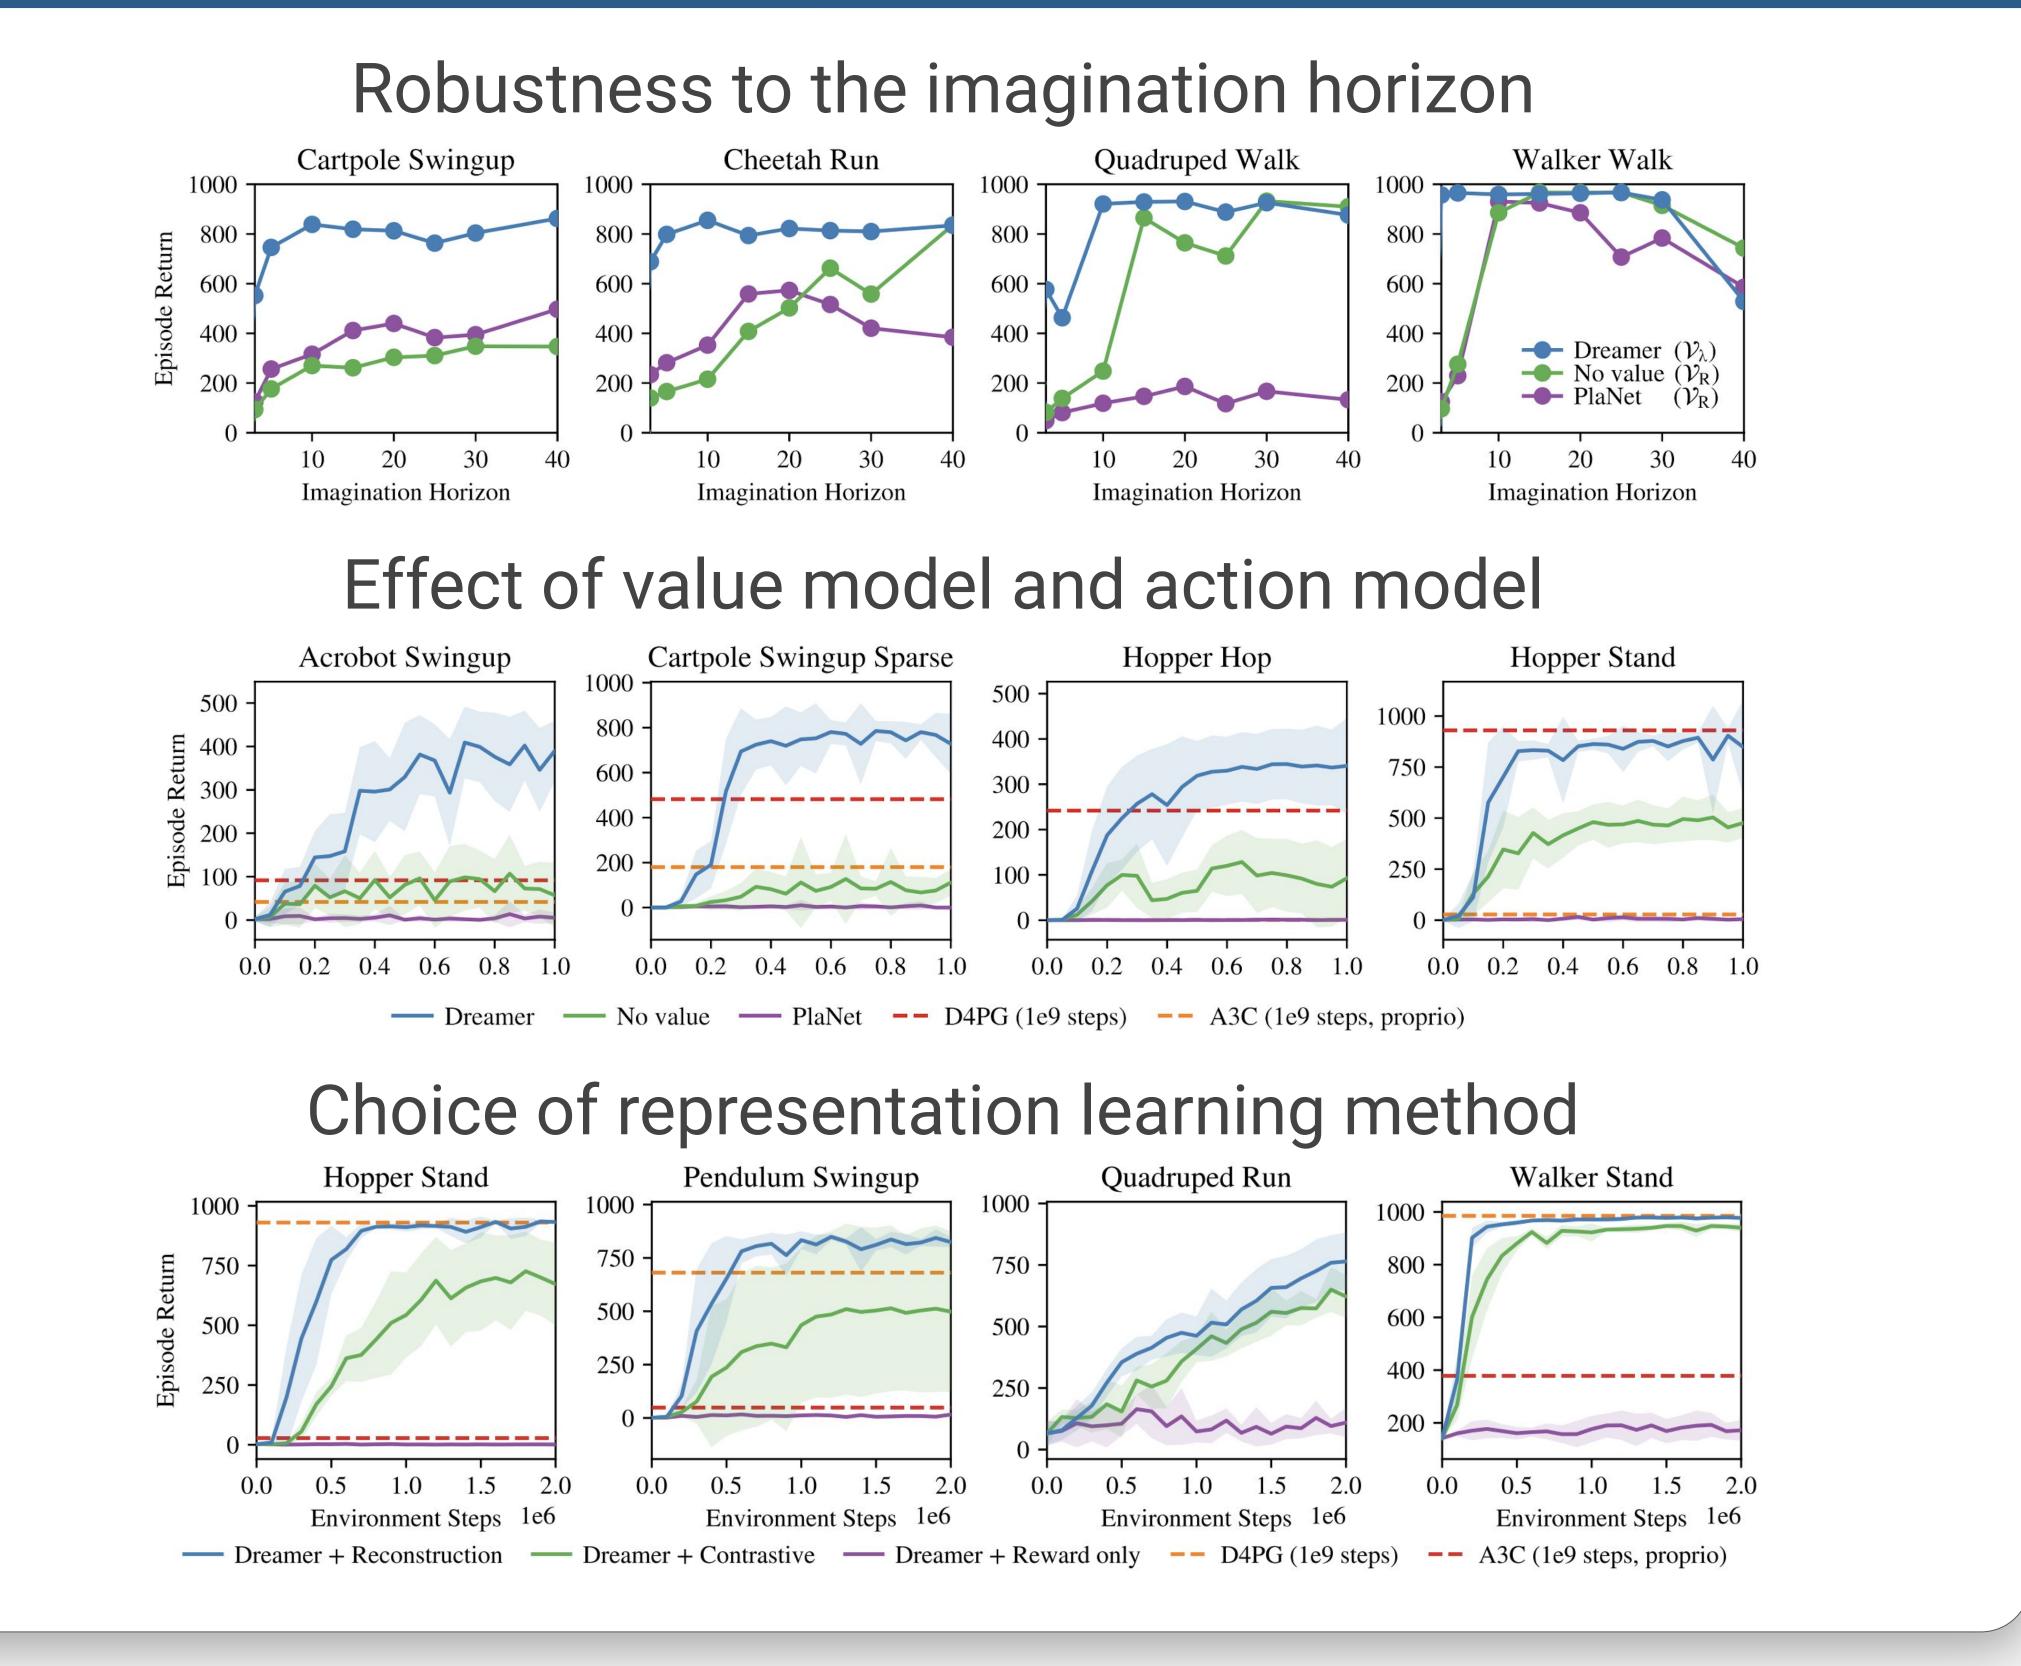
World models are an explicit way to summarize an agent's experience and represent its knowledge about the world.

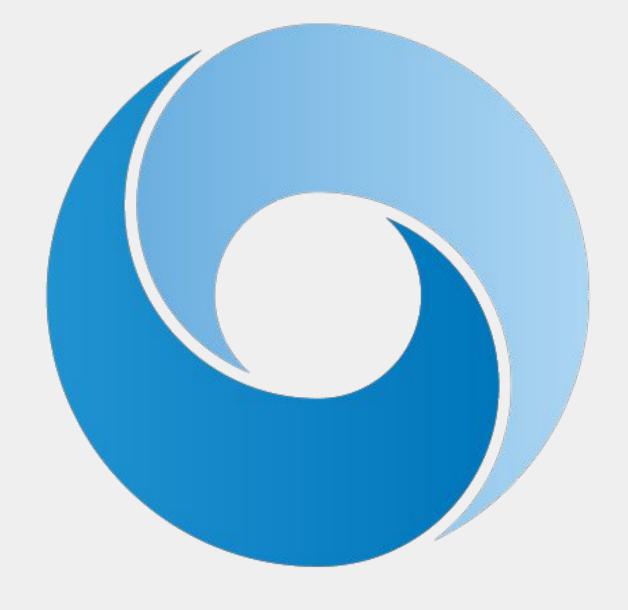
When inputs are large, predicting forward in a compact latent space lets us imagine thousands of trajectories in parallel.

Dreamer is compatible with any representation objective, e.g. reconstruction, contrastive estimation, or reward prediction.



6 Ablation Studies





Project website with videos: danijar.com/dreamer