

Dream-MPC: Gradient-Based Model Predictive Control with Latent Imagination

[Jonathan Spieler](#) and Sven Behnke

Autonomous Intelligent Systems, University of Bonn

International Conference on Machine Learning (ICML) 2026

Why Gradient-Based MPC?

- Model Predictive Control (MPC) allows for planning with learned world models
- Gradient-free, sampling-based MPC methods suffer from the curse of dimensionality
- **But** gradient-based methods perform empirically worse




Can we successfully perform gradient-based MPC with world models?

Dream-MPC




Observation
 s_t

Dream-MPC




→ Forward propagation

Dream-MPC




→ Forward propagation

Dream-MPC



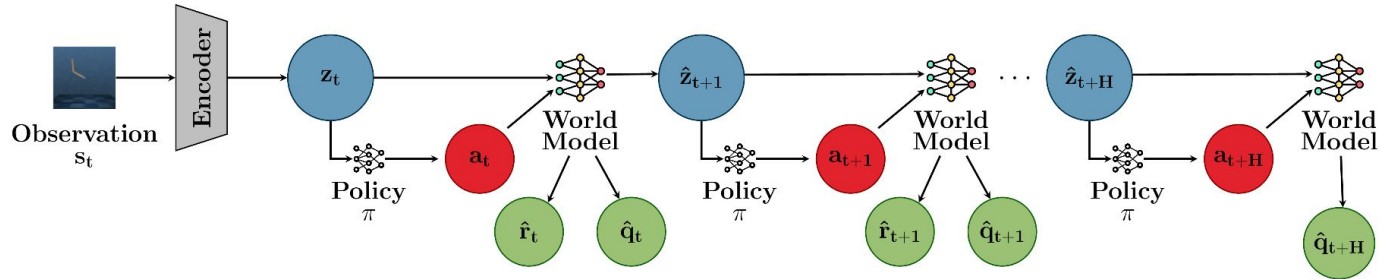
→ Forward propagation

Dream-MPC



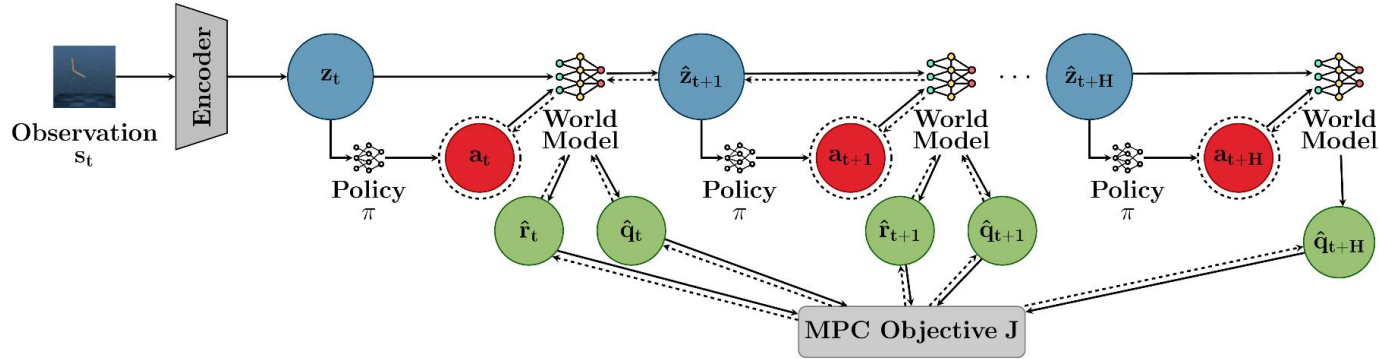
→ Forward propagation

Dream-MPC



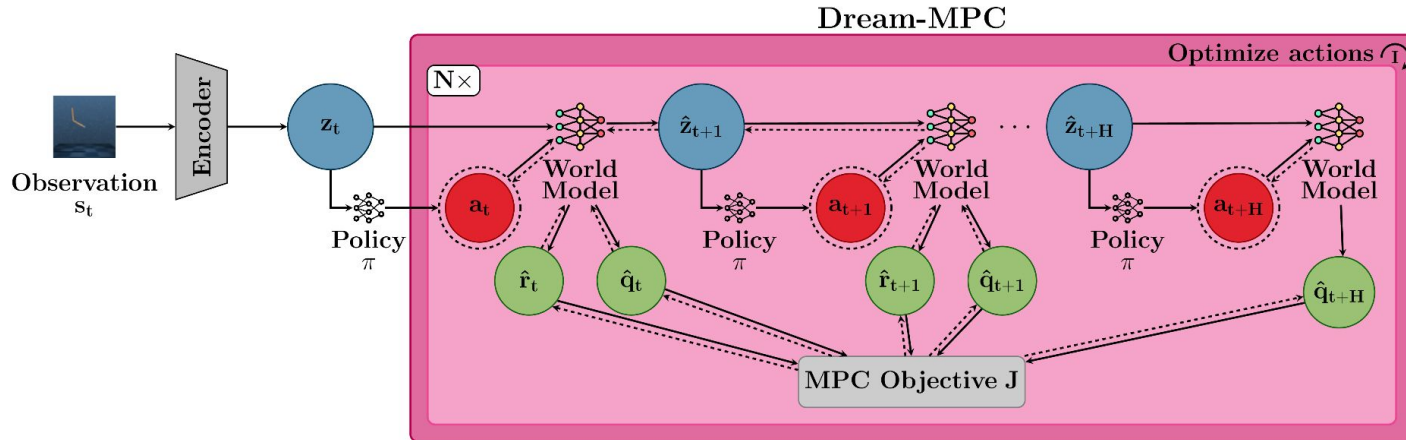
→ Forward propagation

Dream-MPC



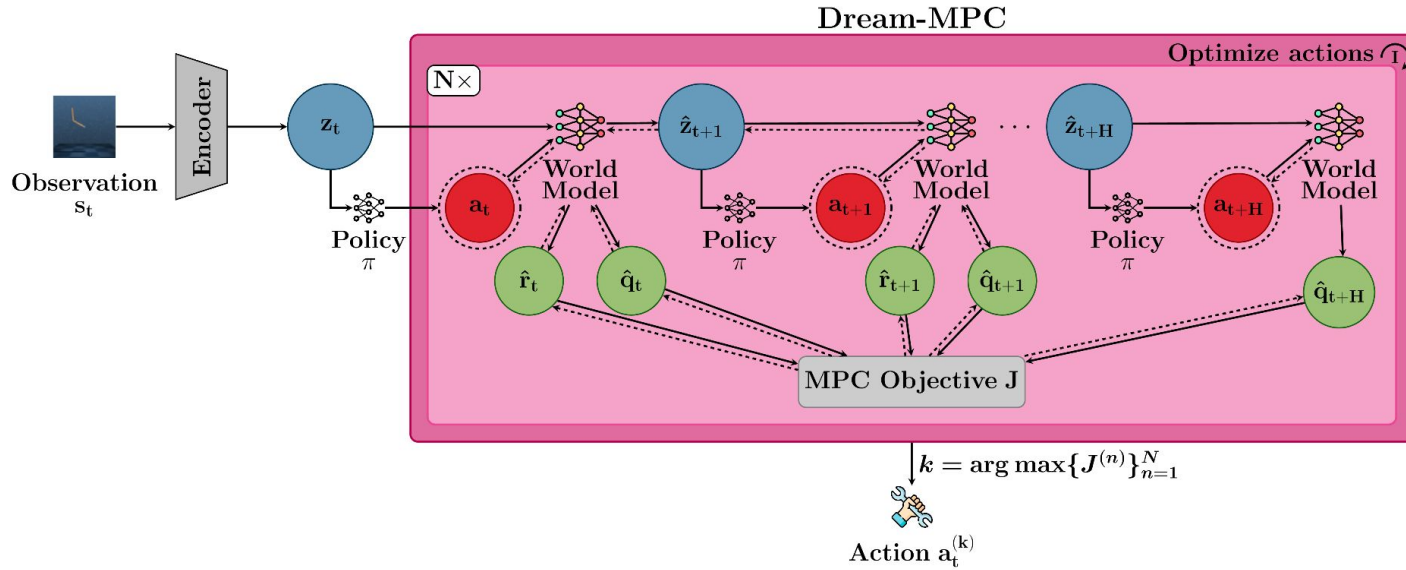
→ Forward propagation ⋯→ Backward propagation

Dream-MPC



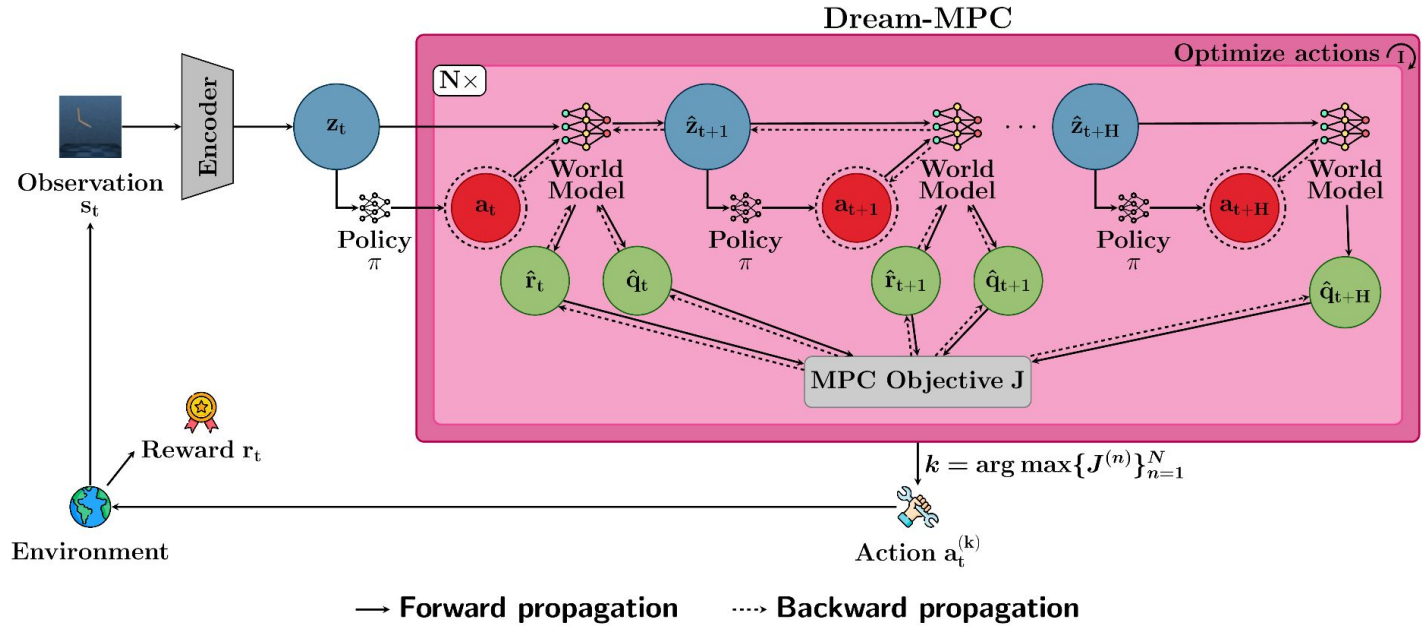
→ Forward propagation ⋯→ Backward propagation

Dream-MPC



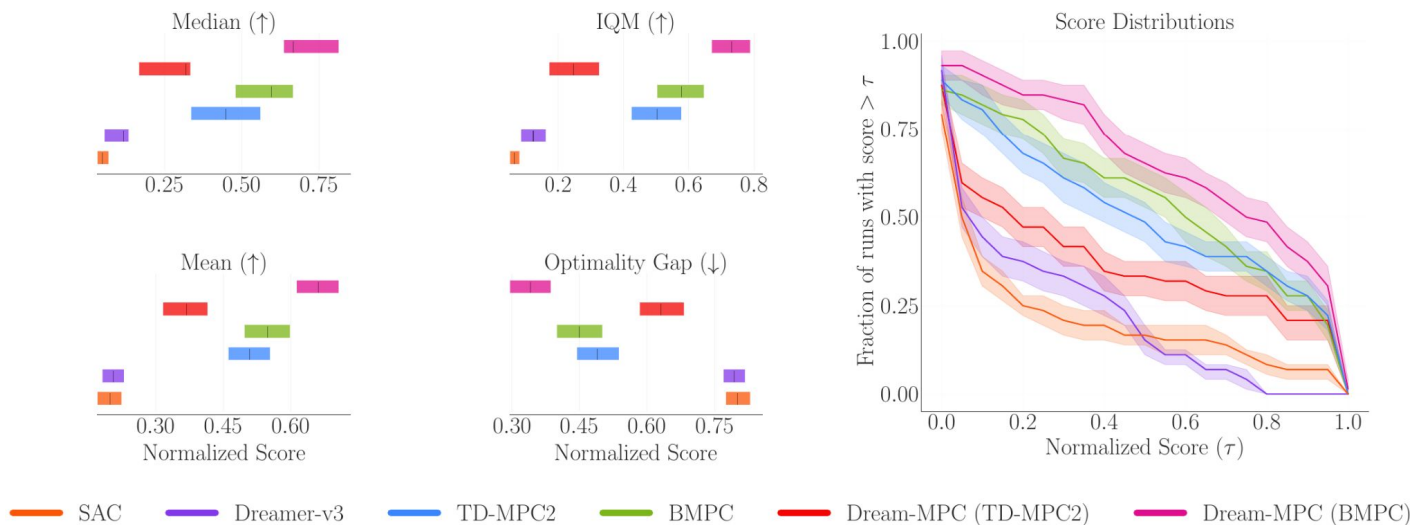
→ Forward propagation ⋯→ Backward propagation

Dream-MPC




How Well Does It Work?

Evaluation of 24 tasks from DM Control, Meta-World & HumanoidBench



What About Other World Models?



Method	Acrobot Swingup	Cheetah Run	Hopper Hop	Walker Run
SAC+AE	7 ± 19	495 ± 100	86 ± 75	453 ± 69
PlaNet	7 ± 18	535 ± 70	1 ± 4	228 ± 149
Dreamer	134 ± 91	751 ± 111	182 ± 43	575 ± 33
Grad-MPC	7 ± 18	438 ± 81	3 ± 5	382 ± 35
Policy+Grad-MPC	144 ± 7	591 ± 131	158 ± 47	556 ± 33
CEM+policy	12 ± 26	674 ± 20	43 ± 42	638 ± 21
Dream-MPC (Dreamer)	147 ± 101	836 ± 49	298 ± 86	632 ± 52

