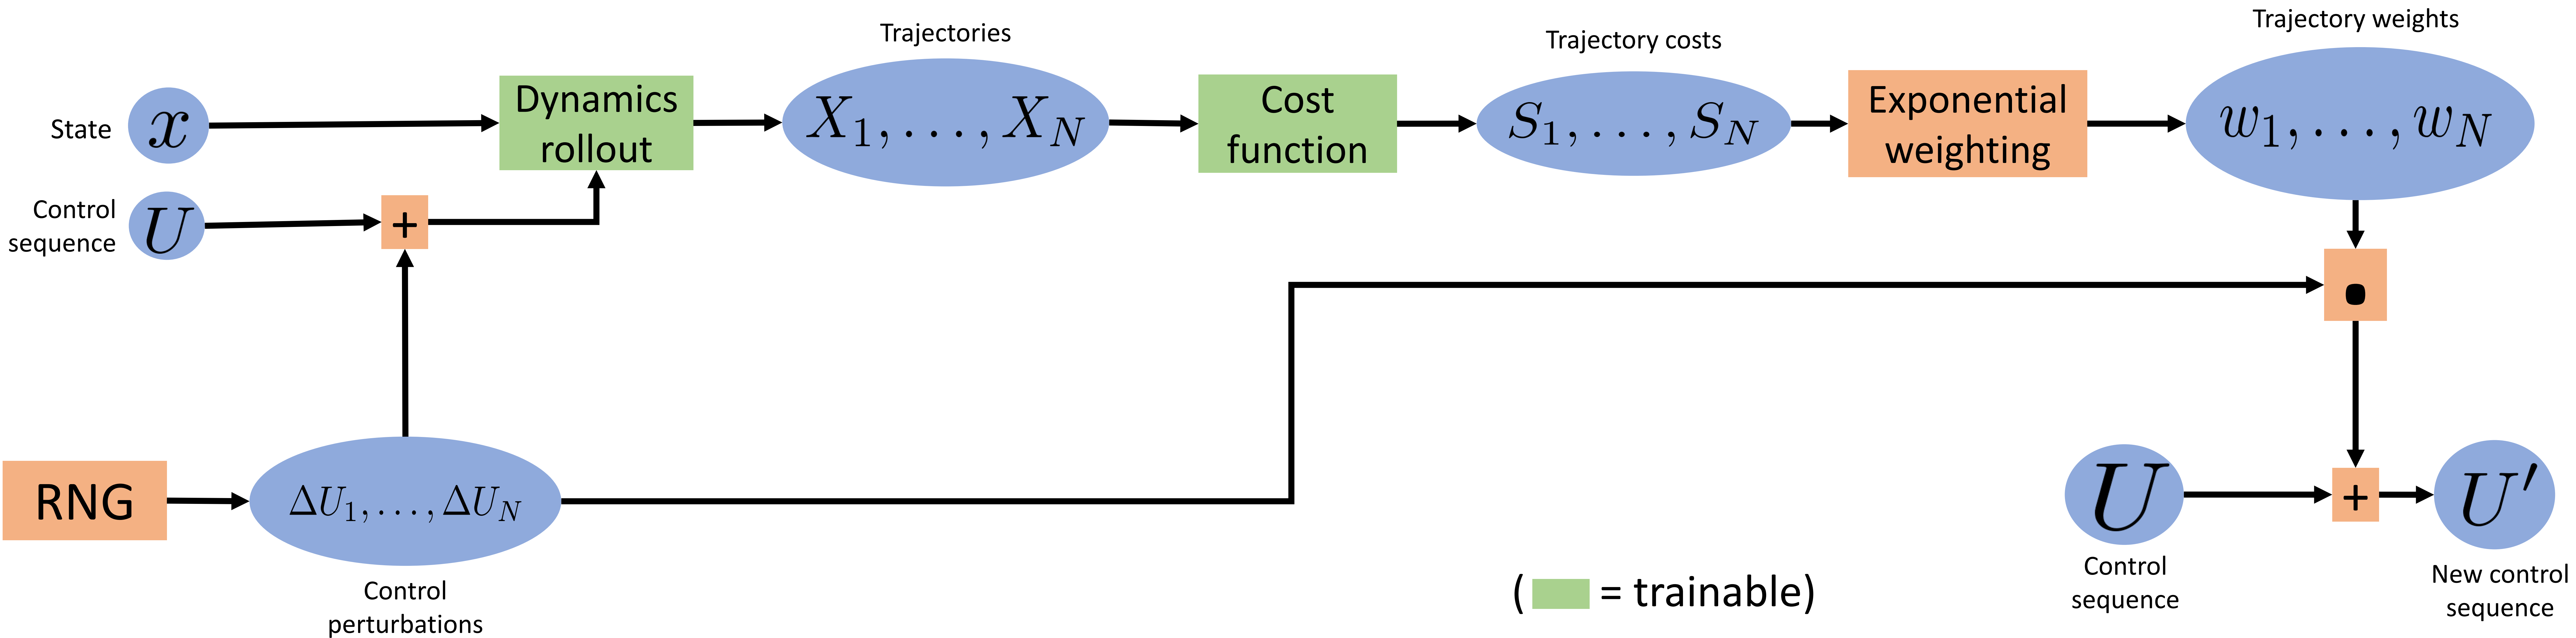


End-to-End Training of an Optimal Control Network


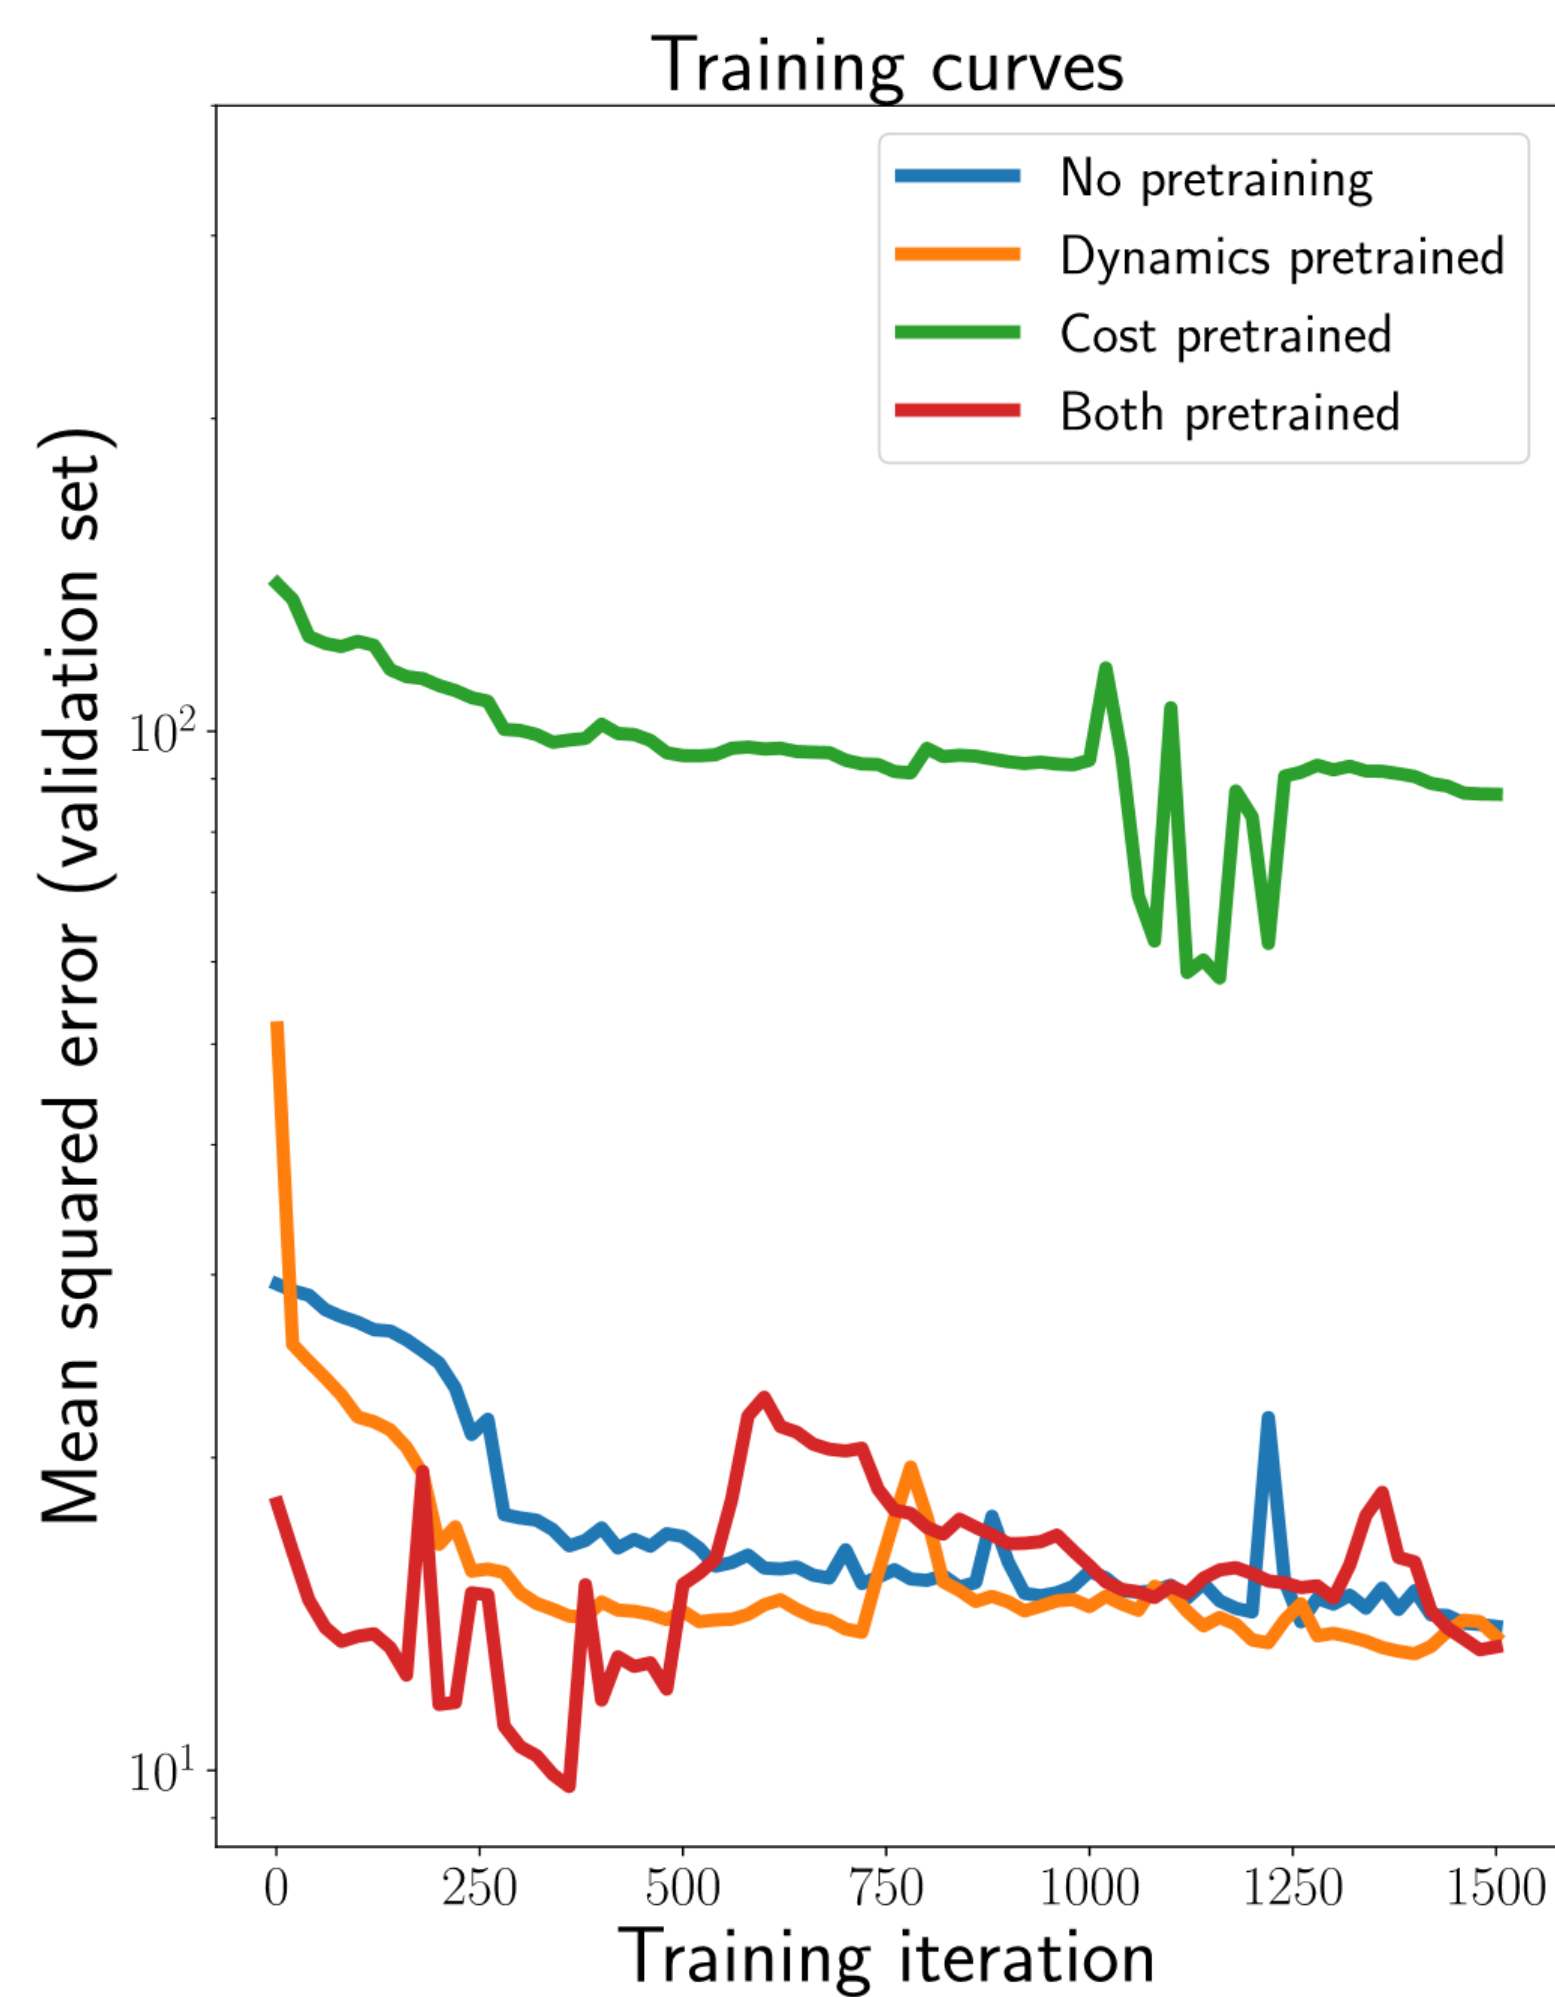
Nolan Wagener (#40)

Path Integral Network (PI-Net): A differentiable policy based on a sampling-based optimal control algorithm



Trained with imitation learning (expert is a PI-Net with access to true dynamics and cost)

Pendulum



Cartpole

