Example-Driven Model-Based Reinforcement Learning for Solving Long-Horizon Visuomotor Tasks

Bohan Wu, Suraj Nair, Li Fei-Fei†, Chelsea Finn† († Equal Advising) Stanford University

Long-Horizon Vision-Based Manipulation

Demands learning a repertoire of visuomotor skills that are:

Robust	**Persistent**
High Success Rates	Closed-loop and Reactive

Key Insight

Use human-provided example images as supervision to learn a repertoire of skills, groundings, and success detectors

1. Provides supervision signal for reward learning in RL
2. Enables skill grounding in long-horizon task planning
3. Encourages closed-loop visuomotor control

Performance Evaluation

(a) Robot Environment (b) Train Objects (c) Validation Objects (d) Test Objects

Successful trials (out of 20) and success rates (%)

	EMBR (Ours)	EMBR w/o $f_{\mathcal{T}}$	EMBR w/o f_{vae}, $f_{\mathcal{T}}$ (Qt-Opt)
Avg. Skill Success	**96.5%**	89.1%	85.9%
Avg. Long-Horizon Task Success	**85.0%**	70.0%	58.3%

* Kalashnikov et al. QT-Opt: Scalable Deep Reinforcement Learning for Vision-Based Robotic Manipulation, CoRL'18

Key Takeaways

1. Q-function is important for low-level skill performance
2. Robustness of replanning w/ model is critical for long-horizon performance

Future Work

1. Reduce the amount of human supervision
2. Expand task scope and handle partial observation
3. Generalization to new environment setup

Example-Driven Model-Based Reinforcement Learning (EMBR)

A framework for learning grounded visuomotor skills sequenced by symbolic planners to complete long-horizon tasks.

Step 1: Learn a repertoire of skills with example-driven model-based reinforcement learning (EMBR)

Step 2: Specify model over skills, using image classifiers for representing pre- and post-conditions

Step 3: Run symbolic planner with the repertoire of visuomotor skills learned in Step 1

Learning a Repertoire of Visuomotor Skills with EMBR

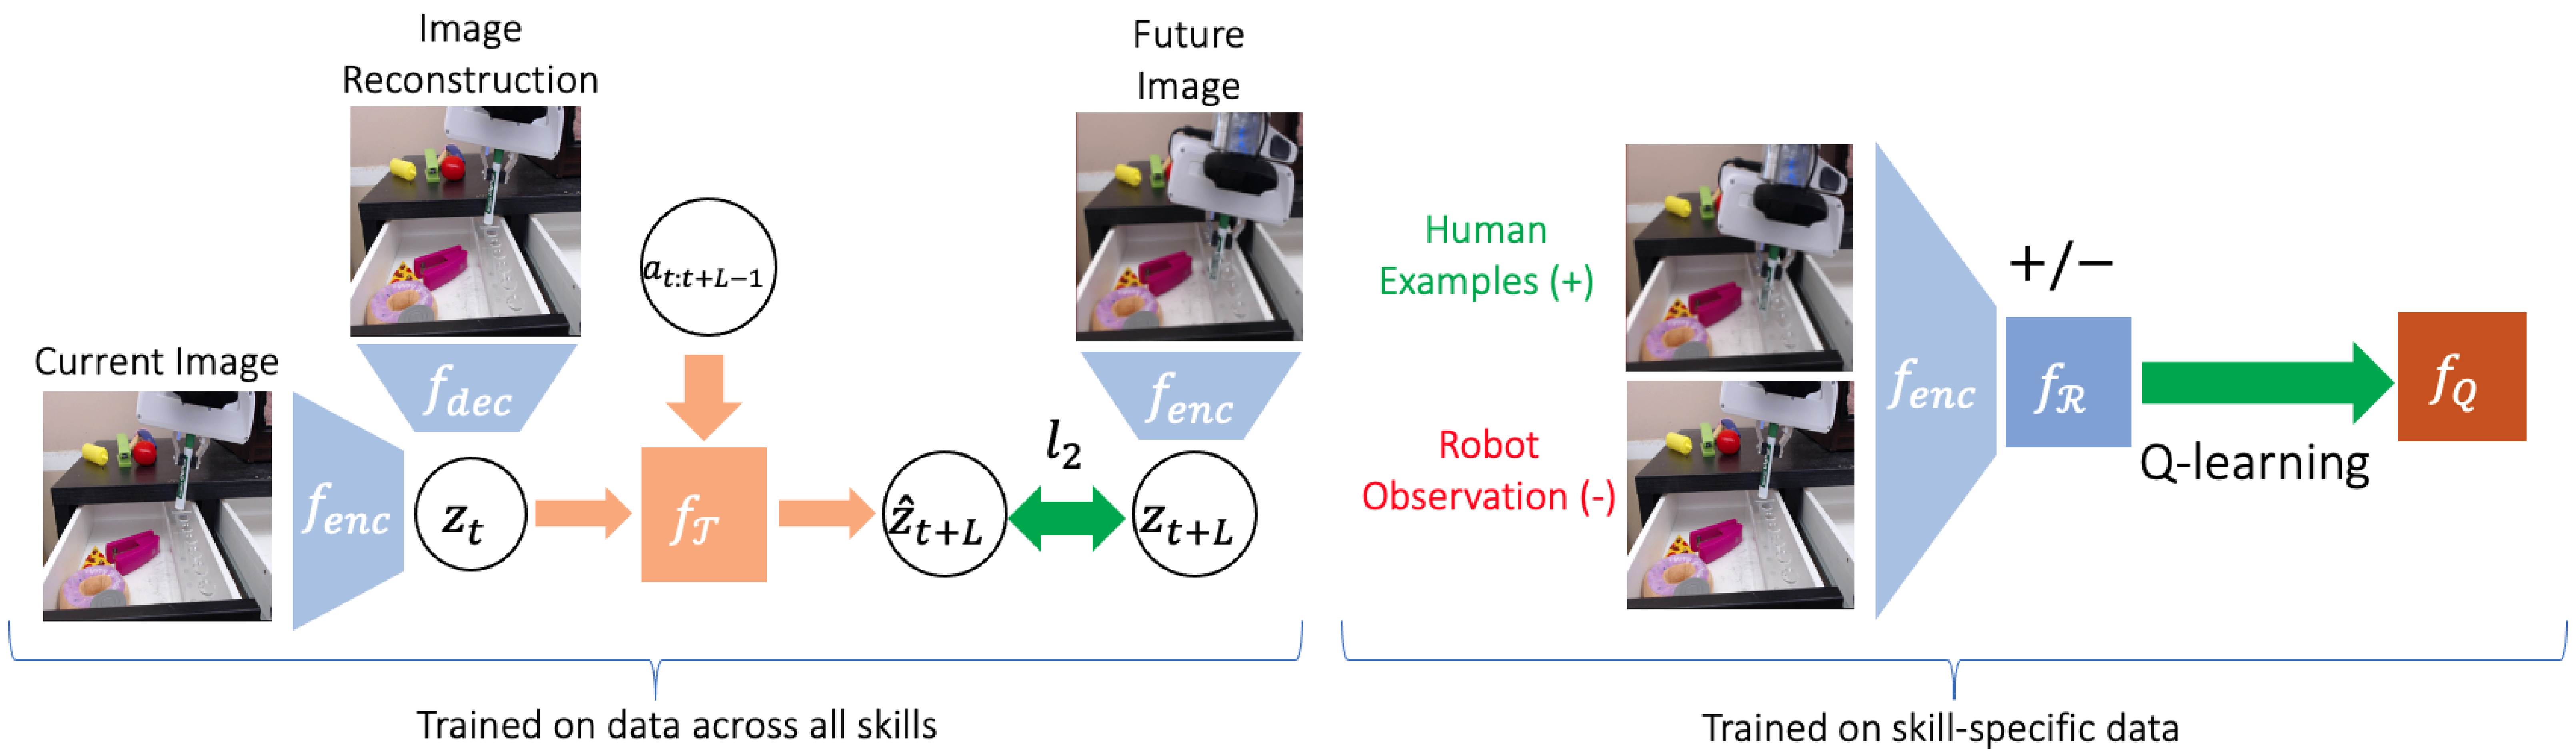

1. Learns a low-dimensional latent space with a VAE and a latent dynamics model
2. Obtains rewards by learning image classifiers
3. Learns Q-functions for model-based control

Real-Robot Experiments: Long-Horizon Tasks with Novel Objects from Raw Image Observations