Diffusion Model-Augmented Behavioral Cloning

Hsiang-Chun Wang* Shang-Fu Chen* Ming-Hao Hsu Chun-Mao Lai Shao-Hua Sun

National Taiwan University NTU Robot Learning Lab

Imitation Learning

Key Idea

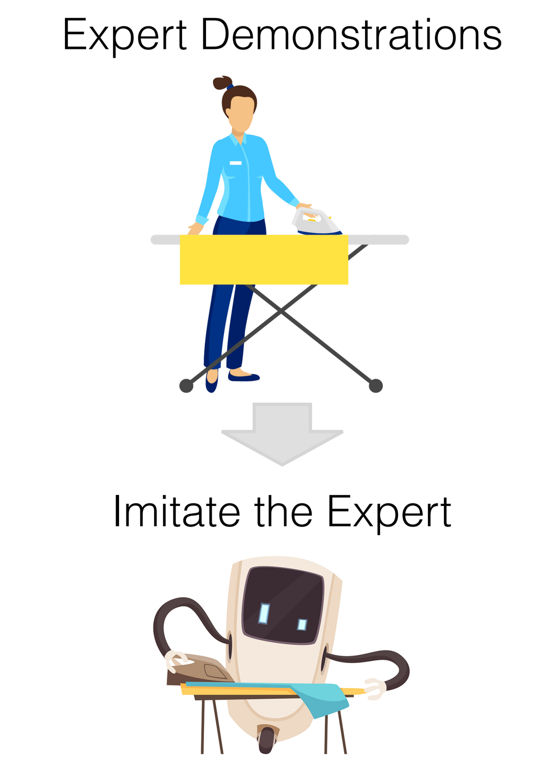

Expert Demonstration

State and action sequences

Modeling Demonstrations

Modeling the conditional probability p(a|s)

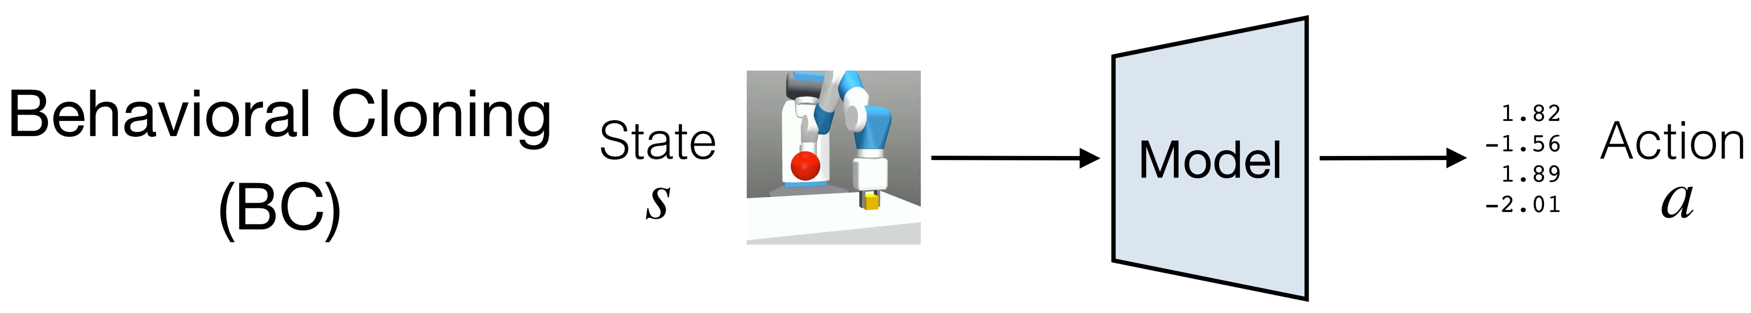

Modeling the joint probability p(s, a)

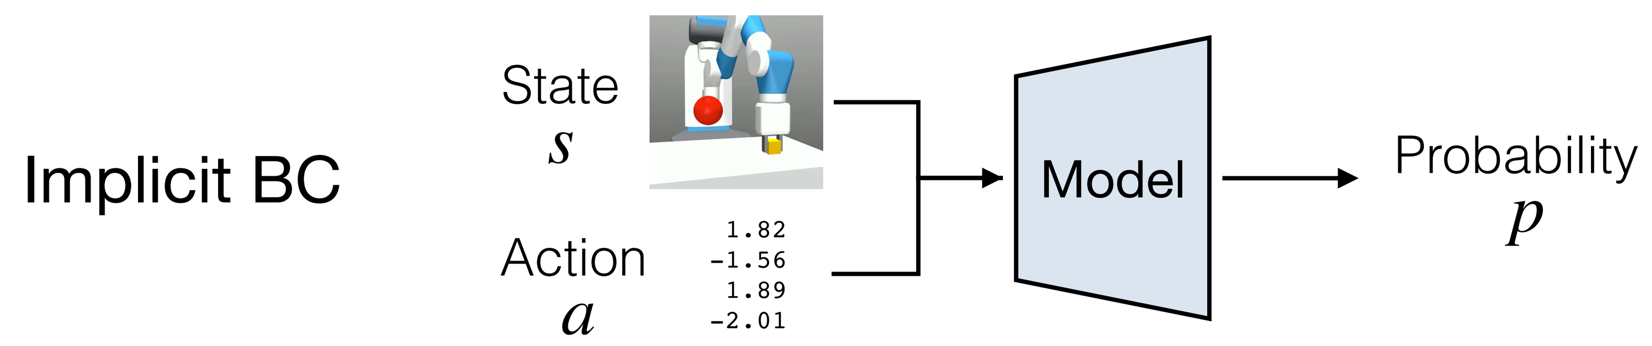

Comparison	Conditional Probability p(a|s)	Joint Probability p(s, a)
Advantages	• Training stability • Inference efficiency	• Better generalization
Disadvantages	• Poor generalization	• Inference inefficiency • Manifold overfitting

- $\pi_{\text{Conditional}}$ **generalizes poorly**
- π_{Joint} suffers from **manifold overfitting**

Main Experiments

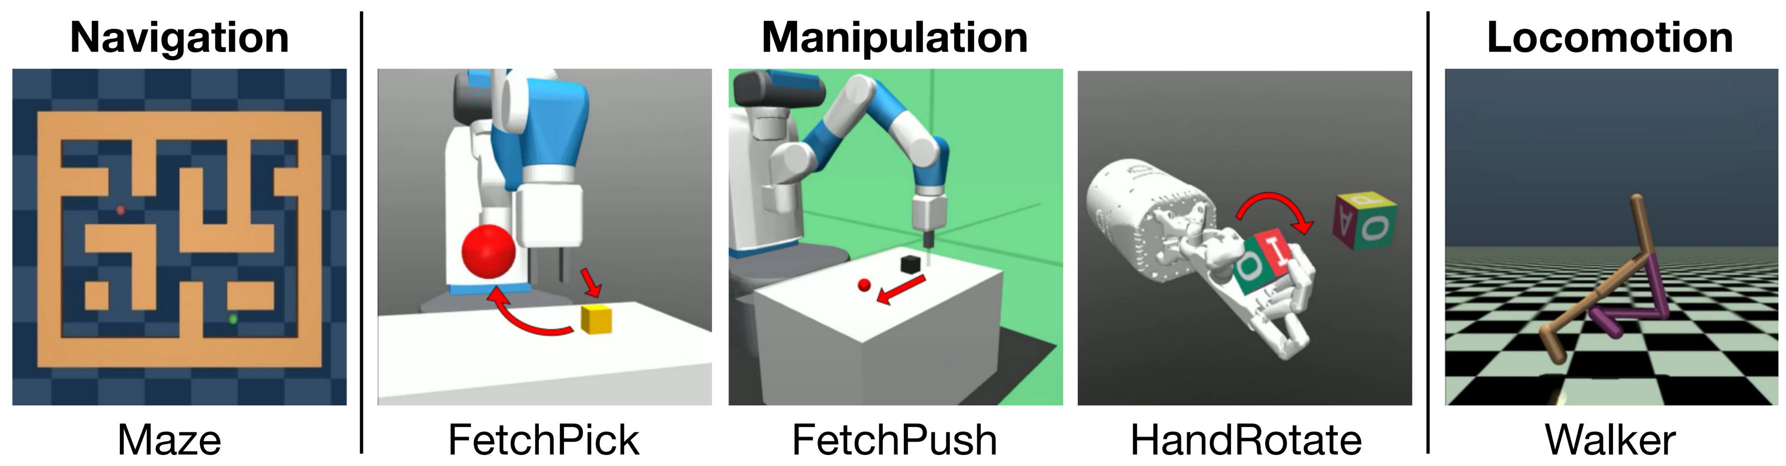

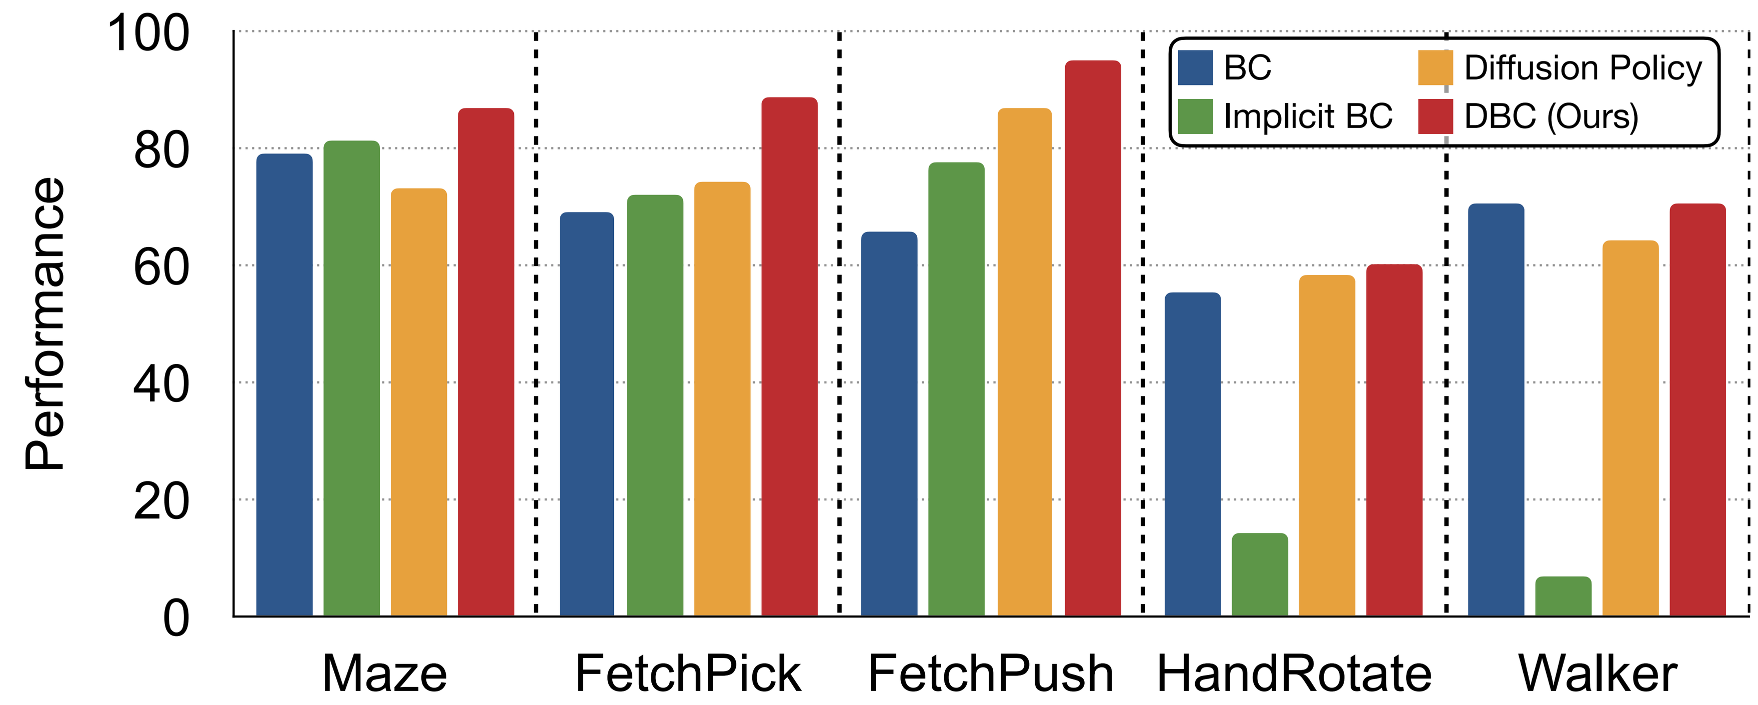

Ablation Studies

Comparing generative models on Maze

- Energy-based model (EBM), variational autoencoder (VAE), generative adversarial network (GAN), diffusion model (DM)

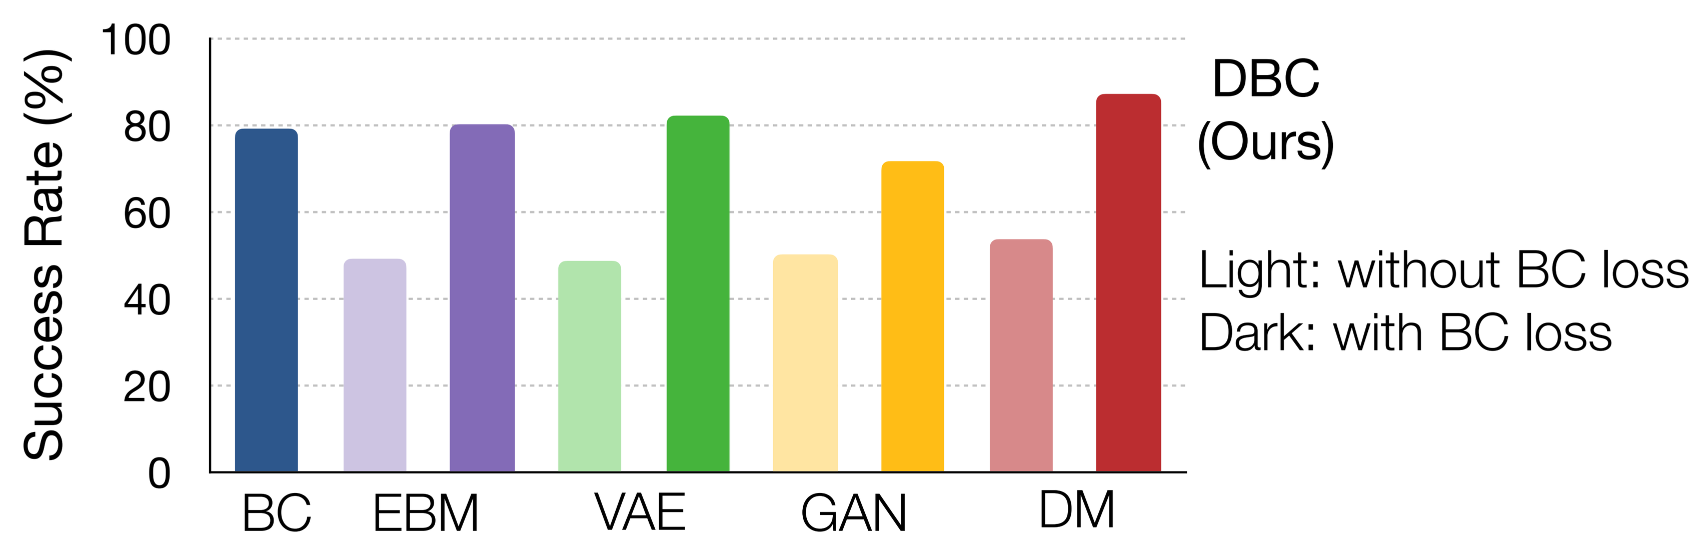

Evaluating generalization performance on Fetch

- Varying the noise added to initial states and goal locations

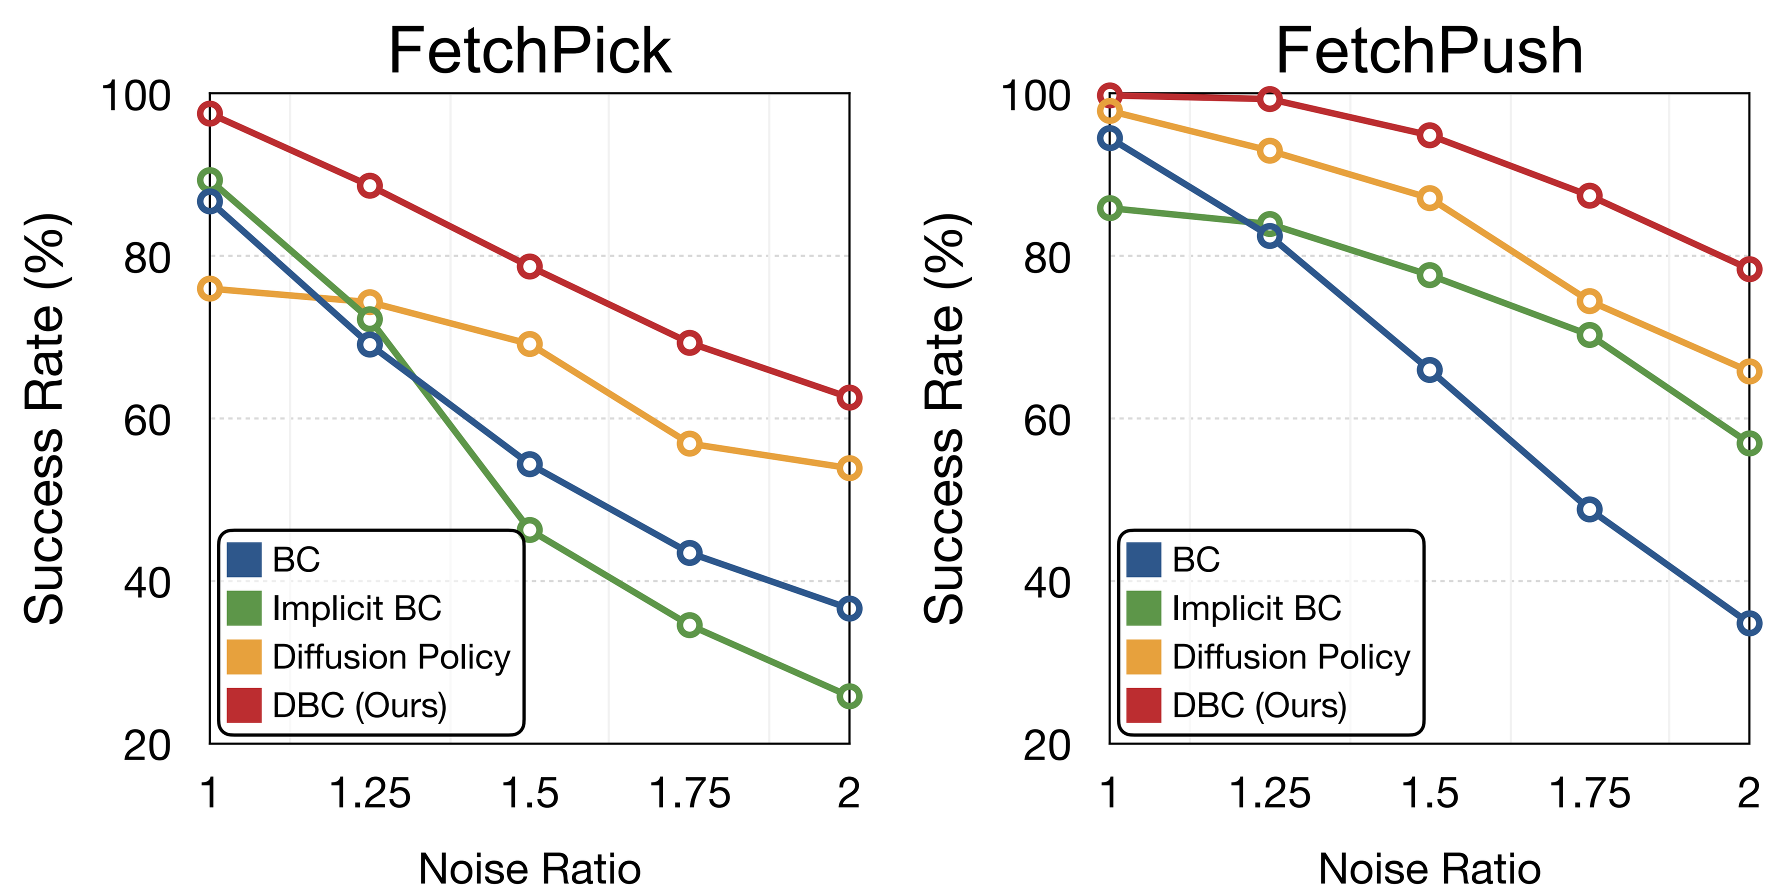

Our Approach: Diffusion Model-Augmented Behavioral Cloning (DBC)

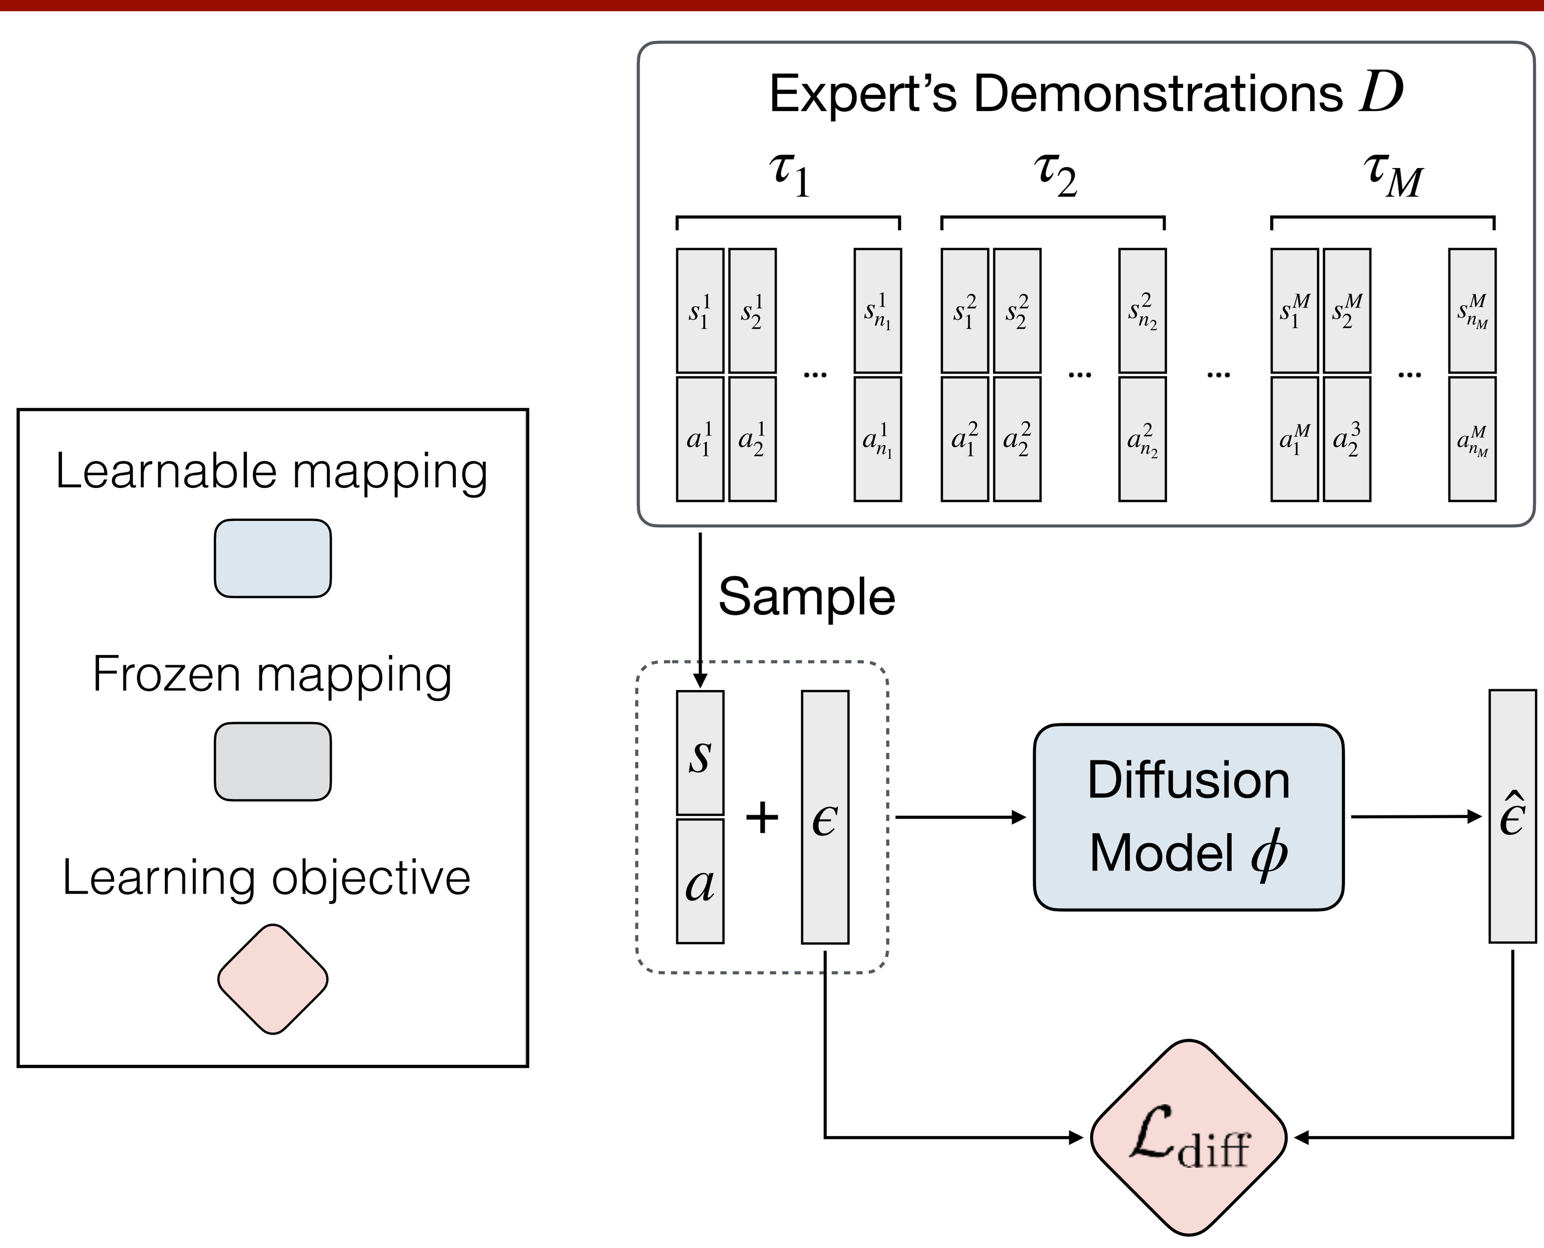

Stage 1: Learning a **Diffusion Model**

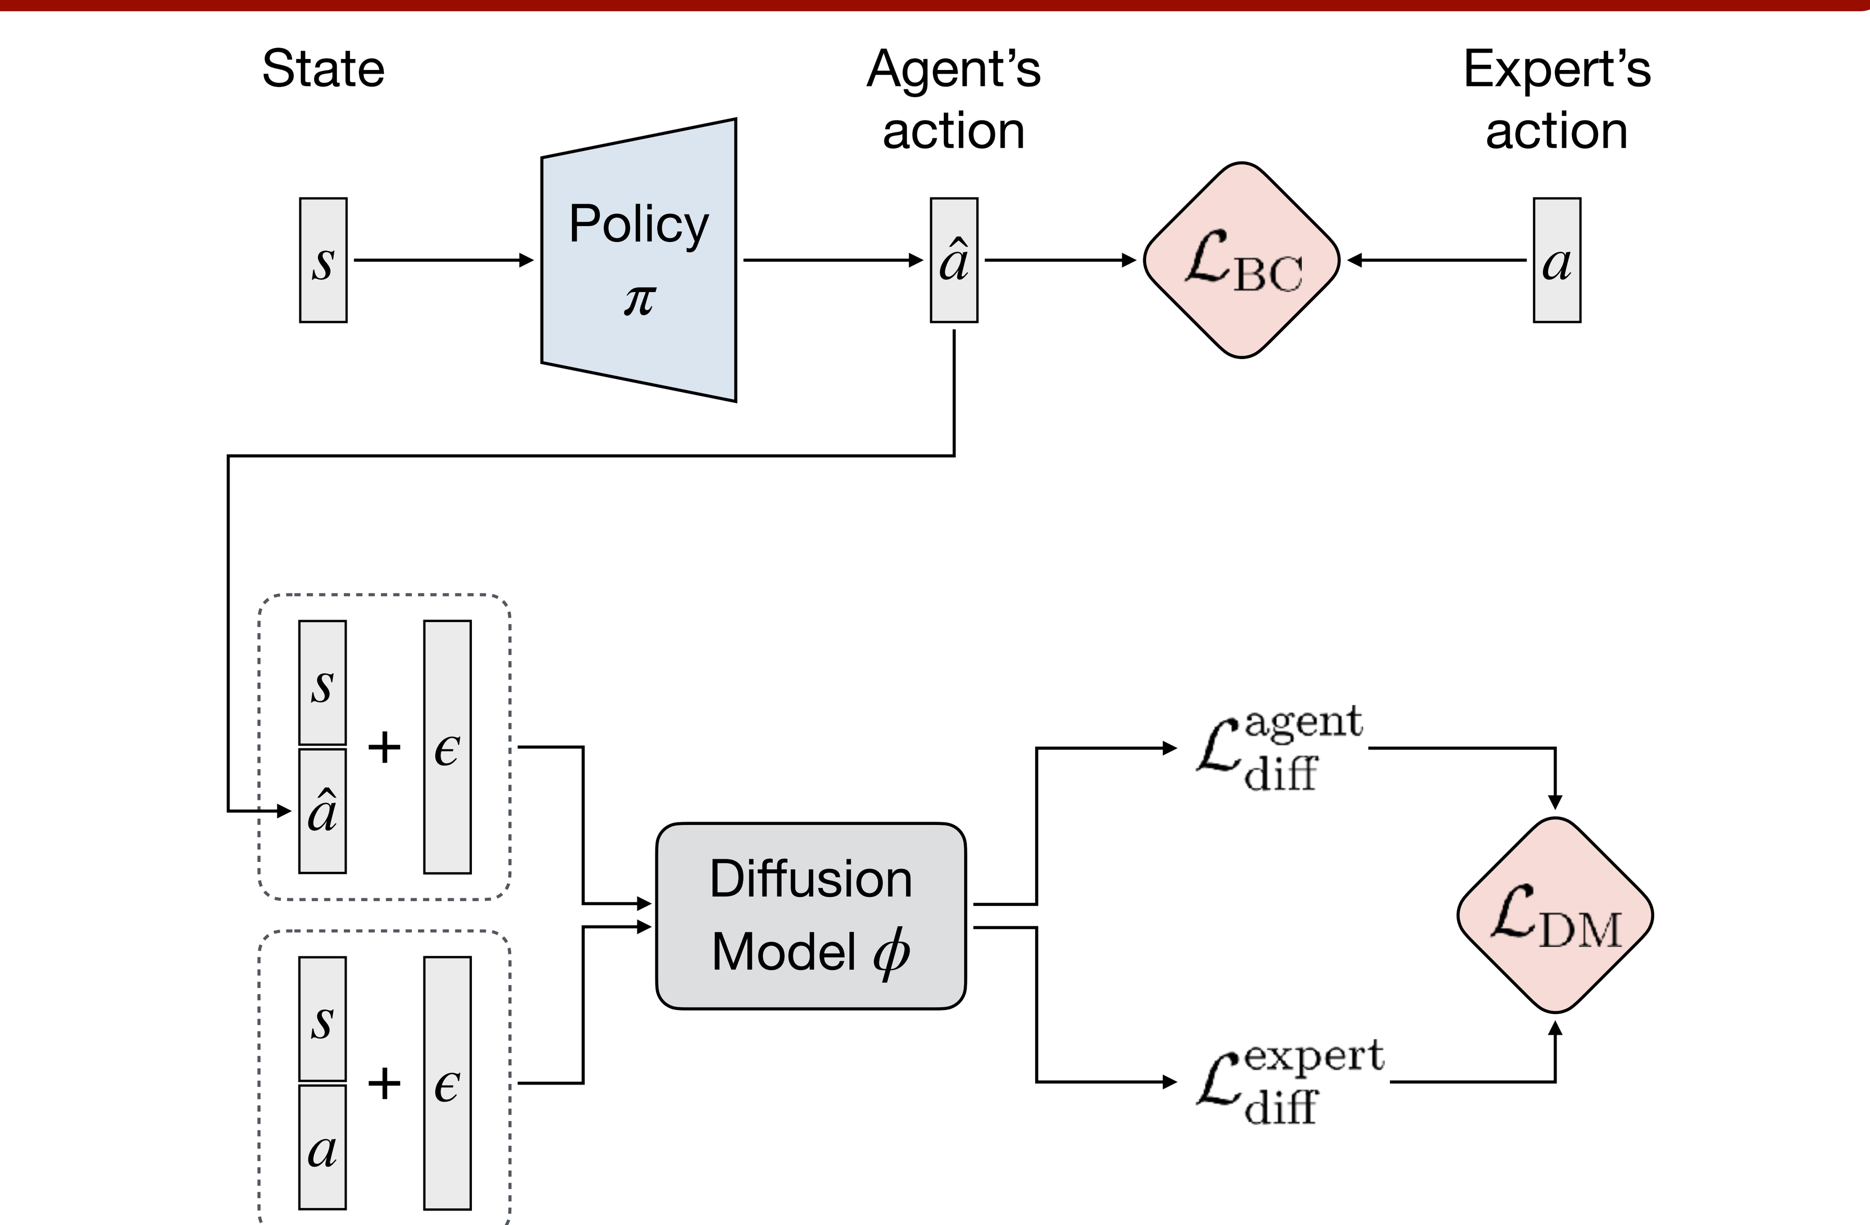

Stage 2: Learning a **Policy** with the **Learned Diffusion Model**