Table 1: Performance of multi-task policies learned using BAKU on LIBERO-90 and Meta-World. We report the mean and standard deviation for each variant across 3 seeds.

Method	LIBERO-90 (90 tasks)	Meta-World (30 tasks)
RT-1	0.14 ± 0.02	0.64 ± 0.01
MTACT	0.55 ± 0.01	0.12 ± 0.01
BAKU **(Ours)**	$\mathbf{0.89 \pm 0.01}$	$\mathbf{0.81 \pm 0.02}$

Table 2: Performance of BAKU with different action heads on LIBERO-90 and Meta-World. We report the mean and standard deviation for each variant across 3 seeds.

Action Head	LIBERO-90	Meta-World
MLP	0.89 ± 0.01	$\mathbf{0.81 \pm 0.02}$
GMM	0.83 ± 0.02	0.64 ± 0.02
BeT	0.88 ± 0.01	0.77 ± 0.01
VQ-BeT	$\mathbf{0.9 \pm 0.01}$	0.78 ± 0.005
Diffusion	0.88 ± 0.01	0.64 ± 0.01

Table 3: Performance of BAKU on long-horizon tasks on LIBERO-10 and a real xArm robot.

Method	LIBERO-10 (10 tasks)	Real Robot (5 tasks)
MT-ACT	0.68	0.64
BAKU **(Ours)**	**0.86**	**0.84**

Table 4: Data efficiency analysis on the LIBERO-90 benchmark.

# Demos	RT-1	MT-ACT	BAKU
5	0	0.31	**0.58**
10	0.01	0.48	**0.71**
25	0.04	0.49	**0.83**
50	0.16	0.54	**0.9**

Table 5: Data efficiency analysis on the Meta-World benchmark.

# Demos	RT-1	MT-ACT	BAKU
5	0.40	0.07	**0.59**
10	0.49	0.10	**0.67**
25	0.62	0.11	**0.76**
35	0.65	0.13	**0.79**

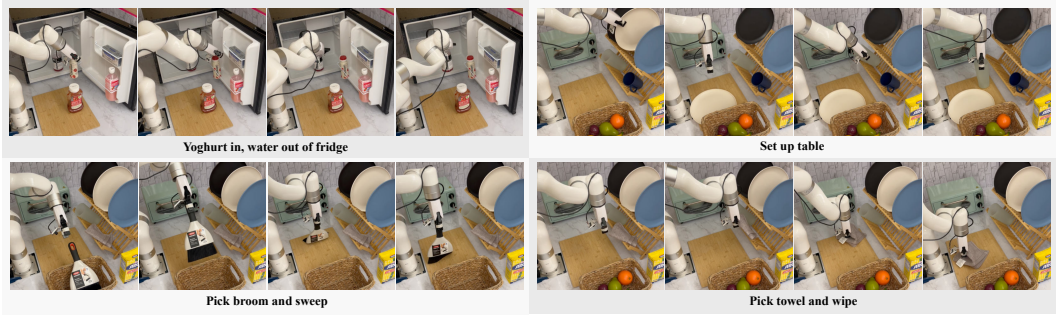

Figure 1: Real-world policy rollouts showing BAKU's capability in long-horizon tasks.