

387 **A Experimental Details**

388 In order to generate the results presented in Table 2 Table 3 and Table 4, we conducted a hyperparam-
389 eter search and selected the best results from the final evaluations for each dataset. Our algorithm
390 was implemented using JAX for the D4RL benchmark. For V-D4RL, we implement our approach
391 using PyTorch adopting the TD3+BC implementation from Clean Offline RL (Tarasov et al., 2022).
392 The experiments were conducted on V100 and A100 GPUs.

393 **Gym-MuJoCo and Adroit tasks.** Our study utilized the latest version of the datasets – v2 for
394 Gym-MuJoCo and v1 for Adroit. The agents were trained for one million steps and evaluated over
395 ten episodes.

396 For ReBRAC, we fine-tuned the β_1 parameter for the actor, which was selected from
397 0.001, 0.01, 0.05, 0.1. Similarly, the β_2 parameter for the critic was selected from a range of
398 0, 0.001, 0.01, 0.1, 0.5. The selected best parameters for each dataset are reported in Table 9.

399 For TD3+BC here and in the AntMaze domain, we use the same grid used in ReBRAC for actor
400 regularization parameter α and add the default value of 0.4.

401 For IQL here and in the AntMaze domain, we selected β value from a range of 0.5, 1, 3, 6, 10 and
402 IQL τ value from a range of 0.5, 0.7, 0.9, 0.95. We used the implementation from Clean Offline RL
403 (Tarasov et al., 2022) and kept other parameters unchanged.

404 For SAC-RND in Adroit domain we tune β_1 (actor parameter) in the range of 0.5, 1.0, 2.5, 5.0, 10.0
405 and β_2 (critic parameter) in the range of 0.01, 0.1, 1.0, 5.0, 10.0.

406 **AntMaze tasks.** In our work, we utilized v2 of the datasets. It’s worth noting that previous studies
407 have reported results using v0 datasets, which were found to contain numerous issues¹. Each agent
408 was trained for 1 million steps and evaluated over 100 episodes. Following Chen et al. (2022a), we
409 modified the reward function by multiplying it by 100.

410 For ReBRAC, the β_1 (actor) and β_2 (critic) hyperparameters were carefully selected from the respec-
411 tive ranges of 0.0005, 0.001, 0.002, 0.003 and 0, 0.0001, 0.0005, 0.001. In addition, the actor and
412 critic learning rates were optimized from 0.0001, 0.0002, 0.0003, 0.0005 and 0.0003, 0.0005, 0.001,
413 respectively. The optimal hyperparameters for each dataset are presented in Table 9.

414 We also modified the γ value for ReBRAC when addressing these tasks, driven by the following
415 motivation. The length of the episodes in AntMaze can be as long as 1000 steps, while the reward is
416 sparse and can only be obtained at the end of the episode. As a result, the discount for the reward with
417 the default γ can be as low as $0.99^{1000} = 4 \cdot 10^{-5}$, which is extremely low for signal propagation,
418 even when multiplying the reward by 100. By increasing γ to 0.999, the minimum discount value
419 becomes $0.999^{1000} = 0.36$, which is more favorable for signal propagation.

420 **V-D4RL.** We used single-task datasets without distraction with a resolution of 84×84 pixels. For
421 ReBRAC β_1 (actor) parameter was selected from the range of $\{0.03, 0.1, 0.3, 1.0\}$ and β_2 (critic)
422 parameter from the range of $\{0.0, 0.001, 0.005, 0.01, 0.1\}$.

¹<https://github.com/Farama-Foundation/D4RL/issues/77>

423 **B Hyperparameters**

424 **B.1 ReBRAC**

Table 8: ReBRAC’s general hyperparameters.

Parameter	Value
optimizer	Adam Kingma & Ba (2014)
batch size	1024 on Gym-MuJoCo, 256 on other
learning rate (all networks)	1e-3 on Gym-MuJoCo, 3e-4 on Adroit and V-D4RL, 1e-4 on Antmaze
tau (τ)	5e-3
hidden dim (all networks)	256
num hidden layers (all networks)	3
gamma (γ)	0.999 on AntMaze, 0.99 on other
nonlinearity	ReLU

Table 9: ReBRAC’s best hyperparameters used in D4RL benchmark.

Task Name	β_1 (actor)	β_2 (critic)
halfcheetah-random	0.001	0.1
halfcheetah-medium	0.001	0.01
halfcheetah-expert	0.01	0.01
halfcheetah-medium-expert	0.01	0.1
halfcheetah-medium-replay	0.01	0.001
halfcheetah-full-replay	0.001	0.1
hopper-random	0.001	0.01
hopper-medium	0.01	0.001
hopper-expert	0.1	0.001
hopper-medium-expert	0.1	0.01
hopper-medium-replay	0.05	0.5
hopper-full-replay	0.01	0.01
walker2d-random	0.01	0.0
walker2d-medium	0.05	0.1
walker2d-expert	0.01	0.5
walker2d-medium-expert	0.01	0.01
walker2d-medium-replay	0.05	0.01
walker2d-full-replay	0.01	0.01
antmaze-umaze	0.003	0.002
antmaze-umaze-diverse	0.003	0.001
antmaze-medium-play	0.001	0.0005
antmaze-medium-diverse	0.001	0.0
antmaze-large-play	0.002	0.001
antmaze-large-diverse	0.002	0.002
pen-human	0.1	0.5
pen-cloned	0.05	0.5
pen-expert	0.01	0.01
door-human	0.1	0.1
door-cloned	0.01	0.1
door-expert	0.05	0.01
hammer-human	0.01	0.5
hammer-cloned	0.1	0.5
hammer-expert	0.01	0.01
relocate-human	0.1	0.01
relocate-cloned	0.1	0.01
relocate-expert	0.05	0.01

Table 10: ReBRAC’s best hyperparameters used in V-D4RL benchmark.

Task Name	β_1 (actor)	β_2 (critic)
walker-walk-random	0.03	0.1
walker-walk-medium	0.03	0.005
walker-walk-expert	0.1	0.01
walker-walk-medium-expert	0.3	0.005
walker-walk-medium-replay	0.3	0.01
cheetah-run-random	0.1	0.01
cheetah-run-medium	0.1	0.1
cheetah-run-expert	0.01	0.01
cheetah-run-medium-expert	1.0	0.001
cheetah-run-medium-replay	0.03	0.1
humanoid-walk-random	1.0	0.01
humanoid-walk-medium	1.0	0.005
humanoid-walk-expert	1.0	0.1
humanoid-walk-medium-expert	1.0	0.005
humanoid-walk-medium-replay	1.0	0.001

Table 11: IQL’s best hyperparameters used in D4RL benchmark.

Task Name	β	IQL τ
halfcheetah-random	3.0	0.95
halfcheetah-medium	3.0	0.95
halfcheetah-expert	6.0	0.9
halfcheetah-medium-expert	3.0	0.7
halfcheetah-medium-replay	3.0	0.95
halfcheetah-full-replay	1.0	0.7
hopper-random	1.0	0.95
hopper-medium	3.0	0.7
hopper-expert	3.0	0.5
hopper-medium-expert	6.0	0.7
hopper-medium-replay	6.0	0.7
hopper-full-replay	10.0	0.9
walker2d-random	0.5	0.9
walker2d-medium	6.0	0.5
walker2d-expert	6.0	0.9
walker2d-medium-expert	1.0	0.5
walker2d-medium-replay	0.5	0.7
walker2d-full-replay	1.0	0.7
antmaze-umaze	10.0	0.7
antmaze-umaze-diverse	10.0	0.95
antmaze-medium-play	6.0	0.9
antmaze-medium-diverse	6.0	0.9
antmaze-large-play	10.0	0.9
antmaze-large-diverse	6.0	0.9
pen-human	1.0	0.95
pen-cloned	10.0	0.9
pen-expert	10.0	0.8
door-human	0.5	0.9
door-cloned	6.0	0.7
door-expert	0.5	0.7
hammer-human	3.0	0.9
hammer-cloned	6.0	0.7
hammer-expert	0.5	0.95
relocate-human	1.0	0.95
relocate-cloned	6.0	0.9
relocate-expert	10.0	0.9

Table 12: TD3+BC’s best hyperparameters used in D4RL benchmark.

Task Name	α
halfcheetah-random	0.001
halfcheetah-medium	0.01
halfcheetah-expert	0.4
halfcheetah-medium-expert	0.1
halfcheetah-medium-replay	0.05
halfcheetah-full-replay	0.01
hopper-random	0.4
hopper-medium	0.05
hopper-expert	0.1
hopper-medium-expert	0.1
hopper-medium-replay	0.4
hopper-full-replay	0.01
walker2d-random	0.001
walker2d-medium	0.4
walker2d-expert	0.05
walker2d-medium-expert	0.1
walker2d-medium-replay	0.1
walker2d-full-replay	0.1
antmaze-umaze	0.4
antmaze-umaze-diverse	0.4
antmaze-medium-play	0.003
antmaze-medium-diverse	0.003
antmaze-large-play	0.003
antmaze-large-diverse	0.003
pen-human	0.1
pen-cloned	0.4
pen-expert	0.4
door-human	0.1
door-cloned	0.4
door-expert	0.1
hammer-human	0.4
hammer-cloned	0.4
hammer-expert	0.4
relocate-human	0.1
relocate-cloned	0.1
relocate-expert	0.4

Table 13: SAC-RND’s best hyperparameters used in D4RL Adroit tasks.

Task Name	β_1 (actor)	β_2 (critic)
pen-human	1.0	10.0
pen-cloned	2.5	0.01
pen-expert	10.0	5.0
door-human	5.0	0.01
door-cloned	5.0	1.0
door-expert	10.0	10.0
hammer-human	10.0	0.01
hammer-cloned	1.0	1.0
hammer-expert	2.5	10.0
relocate-human	5.0	0.01
relocate-cloned	5.0	1.0
relocate-expert	10.0	10.0

428 **C Comparison to Ensemble-based Methods**

429 Comparison of ReBRAC with the ensemble-based methods is presented in Table 14, Table 15, and
430 Table 16. We add the following ensemble-based methods: RORL for each domain (Yang et al., 2022),
431 SAC-N/EDAC (An et al., 2021) for the Gym-MuJoCo and Adroit tasks² and MSG (Ghasemipour
432 et al., 2022) for AntMaze tasks³. The mean-wise best results among algorithms are highlighted with
433 **bold**, and the second-best performance is underlined. Our approach, ReBRAC, shows competitive
434 results on the Gym-MuJoCo datasets. On AntMaze tasks, ReBRAC achieves state-of-the-art results
435 among ensemble-free algorithms and a good score compared to ensemble-based algorithms. And on
436 Adroit tasks, our approach outperforms both families of algorithms.

Table 14: ReBRAC evaluation on the Gym domain. We report the final normalized score averaged over 10 unseen training seeds on v2 datasets. CQL, SAC-N and EDAC scores are taken from An et al. (2021). RORL scores are taken from Yang et al. (2022).

Task Name	Ensemble-free				Ensemble-based			
	TD3+BC	IQL	CQL	SAC-RND	SAC-N	EDAC	RORL	ReBRAC, our
halfcheetah-random	<u>30.9</u> ± 0.4	19.5 ± 0.8	31.1 ± 3.5	27.6 ± 2.1	28.0 ± 0.9	28.4 ± 1.0	28.5 ± 0.8	29.5 ± 1.5
halfcheetah-medium	54.7 ± 0.9	50.0 ± 0.2	46.9 ± 0.4	66.4 ± 1.4	<u>67.5</u> ± 1.2	65.9 ± 0.6	66.8 ± 0.7	65.6 ± 1.0
halfcheetah-expert	93.4 ± 0.4	95.5 ± 2.1	97.3 ± 1.1	102.6 ± 4.2	105.2 ± 2.6	106.8 ± 3.4	105.2 ± 0.7	105.9 ± 1.7
halfcheetah-medium-expert	89.1 ± 5.6	92.7 ± 2.8	95.0 ± 1.4	108.1 ± 1.5	107.1 ± 2.0	106.3 ± 1.9	<u>107.8</u> ± 1.1	101.1 ± 5.2
halfcheetah-medium-replay	45.0 ± 1.1	42.1 ± 3.6	45.3 ± 0.3	51.2 ± 3.2	63.9 ± 0.8	61.3 ± 1.9	<u>61.9</u> ± 1.5	51.0 ± 0.8
halfcheetah-full-replay	75.0 ± 2.5	75.0 ± 0.7	76.9 ± 0.9	81.2 ± 1.3	<u>84.5</u> ± 1.2	84.6 ± 0.9	-	82.1 ± 1.1
hopper-random	8.5 ± 0.6	10.1 ± 5.9	5.3 ± 0.6	19.6 ± 12.4	<u>31.3</u> ± 0.0	25.3 ± 10.4	31.4 ± 0.1	8.1 ± 2.4
hopper-medium	60.9 ± 7.6	65.2 ± 4.2	61.9 ± 6.4	91.1 ± 10.1	100.3 ± 0.3	<u>101.6</u> ± 0.6	104.8 ± 0.1	102.0 ± 1.0
hopper-expert	109.6 ± 3.7	108.8 ± 3.1	106.5 ± 9.1	109.8 ± 0.5	110.3 ± 0.3	110.1 ± 0.1	112.8 ± 0.2	100.1 ± 8.3
hopper-medium-expert	87.8 ± 10.5	85.5 ± 29.7	96.9 ± 15.1	109.8 ± 0.6	110.1 ± 0.3	110.7 ± 0.1	112.7 ± 0.2	107.0 ± 6.4
hopper-medium-replay	55.1 ± 31.7	89.6 ± 13.2	86.3 ± 7.3	97.2 ± 9.0	<u>101.8</u> ± 0.5	101.0 ± 0.5	102.8 ± 0.5	98.1 ± 5.3
hopper-full-replay	97.9 ± 17.5	104.4 ± 10.8	101.9 ± 0.6	107.4 ± 0.8	102.9 ± 0.3	105.4 ± 0.7	-	<u>107.1</u> ± 0.4
walker2d-random	2.0 ± 3.6	11.3 ± 7.0	5.1 ± 1.7	18.7 ± 6.9	21.7 ± 0.0	16.6 ± 7.0	<u>21.4</u> ± 0.2	18.1 ± 4.5
walker2d-medium	77.7 ± 2.9	80.7 ± 3.4	79.5 ± 3.2	<u>92.7</u> ± 1.2	87.9 ± 0.2	92.5 ± 0.8	102.4 ± 1.4	82.5 ± 3.6
walker2d-expert	110.0 ± 0.6	96.9 ± 32.3	109.3 ± 0.1	104.5 ± 22.8	107.4 ± 2.4	<u>115.1</u> ± 1.9	115.4 ± 0.5	112.3 ± 0.2
walker2d-medium-expert	110.4 ± 0.6	112.1 ± 0.5	109.1 ± 0.2	104.6 ± 11.2	<u>116.7</u> ± 0.4	114.7 ± 0.9	121.2 ± 1.5	111.6 ± 0.3
walker2d-medium-replay	68.0 ± 19.2	75.4 ± 9.3	76.8 ± 10.0	<u>89.4</u> ± 3.8	78.7 ± 0.7	87.1 ± 2.4	90.4 ± 0.5	77.3 ± 7.9
walker2d-full-replay	90.3 ± 5.4	97.5 ± 1.4	94.2 ± 1.9	105.3 ± 3.2	94.6 ± 0.5	99.8 ± 0.7	-	<u>102.2</u> ± 1.7
Average w/o full-replay	66.8	70.1	70.1	79.5	82.4	82.9	<u>85.7</u>	78.0
Average	70.3	72.9	73.6	82.6	84.4	85.2	-	81.2

²SAC-N and EDAC score 0 on medium and large AntMaze tasks (Tarasov et al., 2022).³MSG numerical results are not available for Gym-MuJoCo tasks and Adroit tasks were not benchmarked.

Table 15: ReBRAC evaluation on AntMaze domain. We report the final normalized score averaged over 10 unseen training seeds on v2 datasets. CQL scores are taken from Ghasemipour et al. (2022). RORL scores are taken from Yang et al. (2022).

Task Name	Ensemble-free				Ensemble-based		
	TD3+BC	IQL	CQL	SAC-RND	RORL	MSG	ReBRAC, our
antmaze-umaze	66.3 ± 6.2	83.3 ± 4.5	74.0	97.0 ± 1.5	97.7 ± 1.9	97.9 ± 1.3	97.8 ± 1.0
antmaze-umaze-diverse	53.8 ± 8.5	70.6 ± 3.7	84.0	66.0 ± 25.0	90.7 ± 2.9	79.3 ± 3.0	88.3 ± 13.0
antmaze-medium-play	26.5 ± 18.4	64.6 ± 4.9	61.2	38.5 ± 29.4	76.3 ± 2.5	85.9 ± 3.9	84.0 ± 4.2
antmaze-medium-diverse	25.9 ± 15.3	61.7 ± 6.1	53.7	74.7 ± 10.7	69.3 ± 3.3	84.6 ± 5.2	76.3 ± 13.5
antmaze-large-play	0.0 ± 0.0	42.5 ± 6.5	15.8	43.9 ± 29.2	16.3 ± 11.1	64.3 ± 12.7	60.4 ± 26.1
antmaze-large-diverse	0.0 ± 0.0	27.6 ± 7.8	14.9	45.7 ± 28.5	41.0 ± 10.7	71.3 ± 5.3	54.4 ± 25.1
Average	28.7	58.3	50.6	60.9	65.2	80.5	76.8

Table 16: ReBRAC evaluation on Adroit domain. We report the final normalized score averaged over 10 unseen training seeds on v1 datasets. BC, CQL, EDAC and RORL scores are taken from Yang et al. (2022).

Task Name	Ensemble-free				Ensemble-based			
	BC	TD3+BC	IQL	CQL	SAC-RND	RORL	EDAC	ReBRAC, our
pen-human	34.4	81.8 ± 14.9	81.5 ± 17.5	37.5	5.6 ± 5.8	33.7 ± 7.6	51.2 ± 8.6	103.5 ± 14.1
pen-cloned	56.9	61.4 ± 19.3	77.2 ± 17.7	39.2	2.5 ± 6.1	35.7 ± 35.7	68.2 ± 7.3	91.8 ± 21.7
pen-expert	85.1	146.0 ± 7.3	133.6 ± 16.0	107.0	45.4 ± 22.9	130.3 ± 4.2	122.8 ± 14.1	154.1 ± 5.4
door-human	0.5	-0.1 ± 0.0	3.1 ± 2.0	9.9	0.0 ± 0.0	3.7 ± 0.7	10.7 ± 6.8	0.0 ± 0.1
door-cloned	-0.1	0.1 ± 0.6	0.8 ± 1.0	0.4	0.2 ± 0.8	-0.1 ± 0.1	9.6 ± 8.3	1.1 ± 2.6
door-expert	34.9	84.6 ± 44.5	105.3 ± 2.8	101.5	73.6 ± 26.7	104.9 ± 0.9	-0.3 ± 0.1	104.6 ± 2.4
hammer-human	1.5	0.4 ± 0.4	2.5 ± 1.9	4.4	-0.1 ± 0.1	2.3 ± 2.3	0.8 ± 0.4	0.2 ± 0.2
hammer-cloned	0.8	0.8 ± 0.7	1.1 ± 0.5	2.1	0.1 ± 0.4	1.7 ± 1.7	0.3 ± 0.0	6.7 ± 3.7
hammer-expert	125.6	117.0 ± 30.9	129.6 ± 0.5	86.7	24.8 ± 39.4	132.2 ± 0.7	0.2 ± 0.0	133.8 ± 0.7
relocate-human	0.0	-0.2 ± 0.0	0.1 ± 0.1	0.2	0.0 ± 0.0	0.0 ± 0.0	0.1 ± 0.1	0.0 ± 0.0
relocate-cloned	-0.1	-0.1 ± 0.1	0.2 ± 0.4	-0.1	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.9 ± 1.6
relocate-expert	101.3	107.3 ± 1.6	106.5 ± 2.5	95.0	3.4 ± 4.5	47.8 ± 13.5	-0.3 ± 0.0	106.6 ± 3.2
Average w/o expert	11.7	18.0	20.8	11.7	1.0	9.6	17.4	25.5
Average	36.7	49.9	53.4	40.3	12.9	41.0	21.9	58.6

437 D Computational costs

Table 17: Computational costs for algorithms in Table 2.

Algorithm	Number of runs	Approximate hours per run
TD3+BC, tuning	360	0.3
IQL, tuning	1440	1.8
ReBRAC, tuning	1440	0.4
TD3+BC, eval	180	0.2
IQL, eval	180	1.8
SAC-RND, eval	180	1.8
ReBRAC, eval	180	0.4
Sum	3960	4032.0

Table 18: Computational costs for algorithms in Table 3 and Table 15.

Algorithm	Number of runs	Approximate hours per run
TD3+BC, tuning	96	0.5
IQL, tuning	480	2.1
ReBRAC, tuning	384	0.6
TD3+BC, eval	60	0.5
IQL, eval	60	2.0
SAC-RND, eval	60	2.9
MSG, eval	60	5.1
ReBRAC, eval	60	0.4
Sum	1260	1940.4

Table 19: Computational costs for algorithms in Table 4.

Algorithm	Number of runs	Approximate hours per run
TD3+BC, tuning	240	0.3
IQL, tuning	960	1.8
SAC-RND, tuning	1200	1.1
ReBRAC, tuning	960	0.3
TD3+BC, eval	120	0.2
IQL, eval	120	1.9
SAC-RND, eval	120	1.1
ReBRAC, eval	120	0.3
Sum	3840	3828.0

Table 20: Computational costs for algorithms in Table 5.

Algorithm	Number of runs	Approximate hours per run
ReBRAC, tuning	600	10.6
ReBRAC, eval	75	10.5
Sum	675	7147.5

Table 21: Computational costs for algorithms in Table 6 and Figure 2.

Algorithm	Number of runs	Approximate hours per run
ReBRAC, ablations eval	1104	1.4
Sum	1104	1545.6

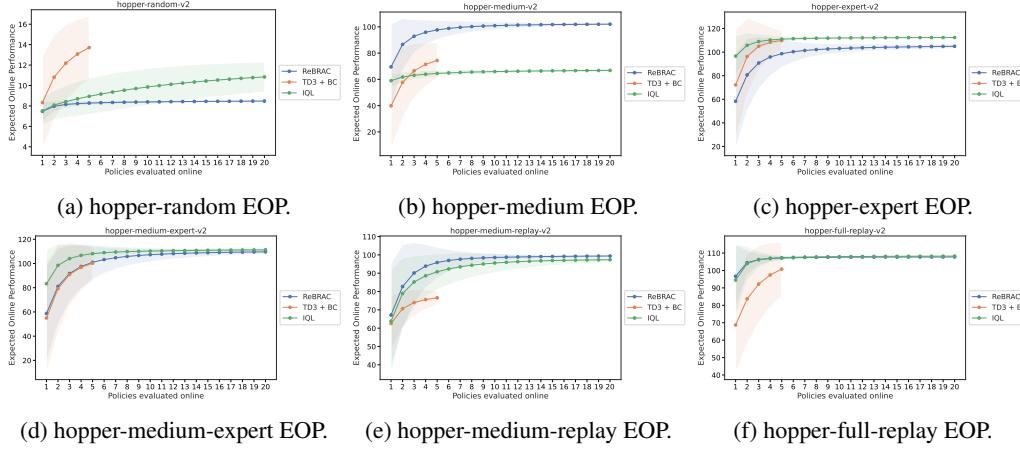


Figure 4: TD3+BC, IQL and ReBRAC visualised Expected Online Performance under uniform policy selection on Hopper tasks.

Table 24: TD3+BC, IQL and ReBRAC Expected Online Performance under uniform policy selection on Walker2d tasks.

Policies	random			medium			expert			medium-expert			medium-replay			full-replay			
	TD3+BC	IQL	ReBRAC	TD3+BC	IQL	ReBRAC	TD3+BC	IQL	ReBRAC	TD3+BC	IQL	ReBRAC	TD3+BC	IQL	ReBRAC	TD3+BC	IQL	ReBRAC	
1	-	6.3 ± 0.7	6.3 ± 0.6	6.3 ± 0.7	6.3 ± 0.6	6.3 ± 0.6	6.3 ± 0.7	6.3 ± 0.6	6.3 ± 0.6	6.3 ± 0.7	6.3 ± 0.6	6.3 ± 0.6	6.3 ± 0.7	6.3 ± 0.6	6.3 ± 0.7	6.3 ± 0.6	6.3 ± 0.6	6.3 ± 0.6	
2	3.7 ± 0.8	7.8 ± 2.8	11.4 ± 6.6	57.0 ± 23.4	75.0 ± 23.4	72.5 ± 23.4	97.9 ± 39.9	110.2 ± 40.4	84.9 ± 32.1	90.9 ± 32.1	108.3 ± 40.4	84.9 ± 32.1	90.9 ± 32.1	96.3 ± 32.2	51.9 ± 24.1	78.5 ± 24.1	71.8 ± 24.1	86.3 ± 24.1	86.3 ± 24.1
3	4.0 ± 0.7	8.8 ± 2.6	13.6 ± 6.3	64.7 ± 17.7	79.0 ± 8.7	79.6 ± 14.9	103.2 ± 26.0	112.7 ± 1.2	109.5 ± 13.0	103.7 ± 20.6	111.2 ± 3.3	109.3 ± 13.4	106.1 ± 13.4	106.1 ± 13.4	50.0 ± 19.0	67.5 ± 21.0	68.8 ± 18.7	87.1 ± 10.5	94.2 ± 10.2
4	4.4 ± 0.6	9.4 ± 2.6	13.8 ± 6.3	65.0 ± 17.7	81.0 ± 8.7	81.6 ± 14.9	103.4 ± 26.0	112.7 ± 1.2	109.5 ± 13.0	103.7 ± 20.6	111.2 ± 3.3	109.3 ± 13.4	106.1 ± 13.4	106.1 ± 13.4	50.4 ± 19.4	74.4 ± 13.1	90.3 ± 6.8	92.4 ± 11.1	99.0 ± 5.7
5	4.3 ± 0.4	9.9 ± 2.3	16.3 ± 5.4	71.3 ± 10.4	81.6 ± 3.9	83.9 ± 6.2	109.1 ± 10.7	114.2 ± 0.8	111.5 ± 3.4	108.8 ± 8.3	111.9 ± 0.9	111.2 ± 3.5	111.2 ± 3.5	111.2 ± 3.5	57.5 ± 10.9	93.4 ± 4.4	93.4 ± 4.4	100.0 ± 1.5	100.0 ± 1.5
6	-	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	10.0 ± 1.0	
7	-	10.5 ± 1.7	17.9 ± 4.6	-	-	-	-	-	-	-	-	-	-	-	79.4 ± 1.1	92.3 ± 3.0	92.3 ± 3.0	96.0 ± 4.8	101.0 ± 3.2
8	-	10.7 ± 1.6	18.4 ± 4.2	-	-	-	-	-	-	-	-	-	-	-	80.0 ± 1.1	94.0 ± 2.8	94.0 ± 2.8	98.0 ± 2.8	101.0 ± 2.5
9	-	10.9 ± 1.4	18.4 ± 3.8	-	-	-	-	-	-	-	-	-	-	-	81.1 ± 3.2	80.9 ± 2.9	97.3 ± 1.9	102.1 ± 2.3	102.1 ± 2.3
10	-	11.0 ± 1.3	19.0 ± 3.5	-	-	-	-	-	-	-	-	-	-	-	81.4 ± 2.8	81.2 ± 2.8	97.5 ± 1.2	102.5 ± 1.8	102.5 ± 1.8
11	-	11.1 ± 1.2	19.2 ± 3.2	-	-	-	-	-	-	-	-	-	-	-	81.6 ± 2.1	81.4 ± 2.1	97.5 ± 1.2	102.5 ± 1.8	102.5 ± 1.8
12	-	11.2 ± 1.2	19.8 ± 3.0	-	-	-	-	-	-	-	-	-	-	-	81.8 ± 1.9	81.6 ± 1.9	97.7 ± 0.9	102.8 ± 1.5	102.8 ± 1.5
13	-	11.3 ± 1.1	19.8 ± 2.8	-	-	-	-	-	-	-	-	-	-	-	81.9 ± 1.7	81.7 ± 1.8	-	97.7 ± 0.9	102.8 ± 1.5
14	-	11.4 ± 1.0	20.4 ± 2.6	-	-	-	-	-	-	-	-	-	-	-	82.0 ± 1.4	81.8 ± 1.4	-	97.8 ± 0.7	102.9 ± 1.2
15	-	11.5 ± 1.0	20.4 ± 2.4	-	-	-	-	-	-	-	-	-	-	-	82.2 ± 1.3	82.0 ± 1.6	-	97.9 ± 0.7	103.0 ± 1.2
16	-	11.5 ± 0.9	20.4 ± 2.2	-	-	-	-	-	-	-	-	-	-	-	82.3 ± 1.3	82.1 ± 1.5	-	97.9 ± 0.6	103.0 ± 0.9
17	-	11.6 ± 0.9	20.7 ± 2.1	-	-	-	-	-	-	-	-	-	-	-	82.4 ± 1.2	82.2 ± 1.4	-	97.9 ± 0.5	103.1 ± 0.8
18	-	11.6 ± 0.8	20.8 ± 1.9	-	-	-	-	-	-	-	-	-	-	-	82.4 ± 1.1	82.2 ± 1.4	-	97.9 ± 0.5	103.2 ± 0.8
19	-	11.7 ± 0.8	20.8 ± 1.8	-	-	-	-	-	-	-	-	-	-	-	82.5 ± 1.0	82.3 ± 1.3	-	97.9 ± 0.5	103.2 ± 0.8
20	-	11.7 ± 0.7	21.0 ± 1.7	-	-	-	-	-	-	-	-	-	-	-	82.5 ± 1.0	82.4 ± 1.3	-	98.0 ± 0.4	103.2 ± 0.6

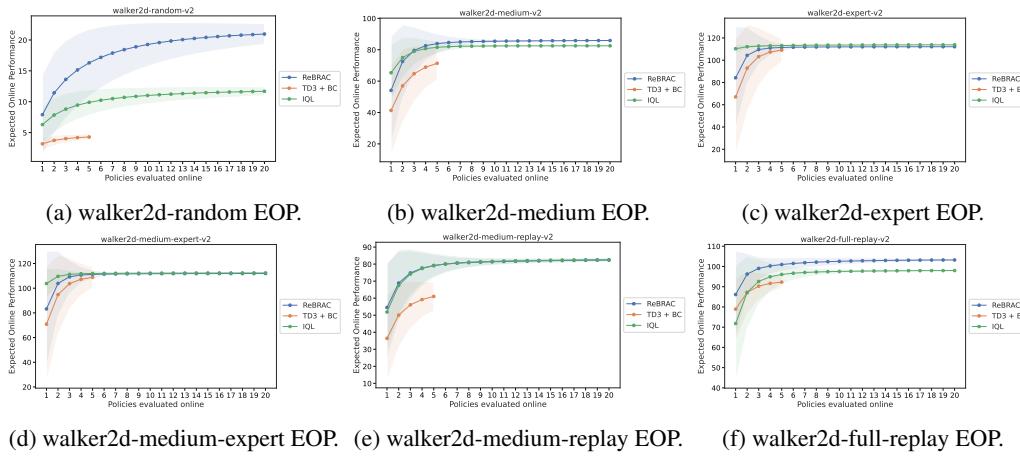


Figure 5: TD3+BC, IQL and ReBRAC visualised Expected Online Performance under uniform policy selection on Walker2d tasks.

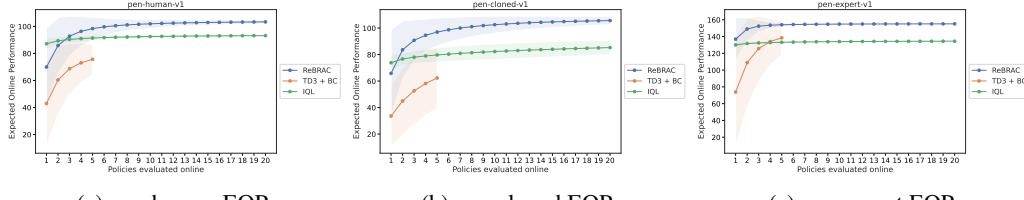


Figure 7: TD3+BC, IQL and ReBRAC visualised Expected Online Performance under uniform policy selection on Pen tasks.

Table 27: TD3+BC, IQL and ReBRAC Expected Online Performance under uniform policy selection on Door tasks.

Policies	human			cloned			expert		
	TD3+BC	IQL	ReBRAC	TD3+BC	IQL	ReBRAC	TD3+BC	IQL	ReBRAC
1	-0.2 ± 0.1	4.4 ± 1.2	-0.1 ± 0.1	-0.1 ± 0.3	1.6 ± 0.8	0.3 ± 0.9	50.7 ± 46.3	102.1 ± 5.7	75.4 ± 43.0
2	-0.1 ± 0.1	5.0 ± 1.0	-0.0 ± 0.1	0.0 ± 0.3	2.0 ± 0.6	0.6 ± 1.2	75.6 ± 39.5	104.7 ± 2.5	96.1 ± 25.0
3	-0.1 ± 0.1	5.4 ± 1.0	-0.0 ± 0.0	0.1 ± 0.3	2.2 ± 0.5	0.8 ± 1.4	88.2 ± 30.4	105.4 ± 1.3	102.3 ± 13.5
4	-0.1 ± 0.1	5.6 ± 0.9	-0.0 ± 0.0	0.2 ± 0.3	2.4 ± 0.4	1.0 ± 1.6	94.9 ± 22.9	105.7 ± 0.8	104.4 ± 7.3
5	-0.1 ± 0.1	5.8 ± 0.9	-0.0 ± 0.0	0.2 ± 0.3	2.5 ± 0.4	1.2 ± 1.7	98.6 ± 17.2	105.8 ± 0.6	105.2 ± 4.1
6	-	5.9 ± 0.9	0.0 ± 0.0	-	2.5 ± 0.3	1.3 ± 1.8	-	105.9 ± 0.5	105.6 ± 2.4
7	-	6.0 ± 0.9	0.0 ± 0.0	-	2.6 ± 0.3	1.5 ± 1.8	-	106.0 ± 0.5	105.8 ± 1.5
8	-	6.1 ± 0.8	0.0 ± 0.0	-	2.6 ± 0.3	1.6 ± 1.9	-	106.1 ± 0.4	105.9 ± 1.0
9	-	6.2 ± 0.8	0.0 ± 0.0	-	2.6 ± 0.2	1.8 ± 1.9	-	106.1 ± 0.4	105.9 ± 0.7
10	-	6.3 ± 0.8	0.0 ± 0.0	-	2.7 ± 0.2	1.9 ± 1.9	-	106.1 ± 0.4	106.0 ± 0.5
11	-	6.4 ± 0.8	0.0 ± 0.0	-	2.7 ± 0.2	2.0 ± 2.0	-	106.2 ± 0.4	106.0 ± 0.4
12	-	6.4 ± 0.7	0.0 ± 0.0	-	2.7 ± 0.2	2.2 ± 2.0	-	106.2 ± 0.3	106.0 ± 0.3
13	-	6.5 ± 0.7	0.0 ± 0.0	-	2.7 ± 0.2	2.3 ± 2.0	-	106.2 ± 0.3	106.0 ± 0.2
14	-	6.5 ± 0.7	0.0 ± 0.0	-	2.7 ± 0.2	2.4 ± 2.0	-	106.2 ± 0.3	106.0 ± 0.2
15	-	6.6 ± 0.7	0.0 ± 0.0	-	2.7 ± 0.2	2.5 ± 2.0	-	106.3 ± 0.3	106.0 ± 0.2
16	-	6.6 ± 0.6	0.0 ± 0.0	-	2.7 ± 0.2	2.6 ± 1.9	-	106.3 ± 0.3	106.1 ± 0.2
17	-	6.6 ± 0.6	0.0 ± 0.0	-	2.8 ± 0.2	2.7 ± 1.9	-	106.3 ± 0.3	106.1 ± 0.2
18	-	6.7 ± 0.6	0.0 ± 0.0	-	2.8 ± 0.2	2.7 ± 1.9	-	106.3 ± 0.3	106.1 ± 0.1
19	-	6.7 ± 0.6	0.0 ± 0.0	-	2.8 ± 0.2	2.8 ± 1.9	-	106.3 ± 0.3	106.1 ± 0.1
20	-	6.7 ± 0.5	0.0 ± 0.0	-	2.8 ± 0.2	2.9 ± 1.9	-	106.3 ± 0.2	106.1 ± 0.1

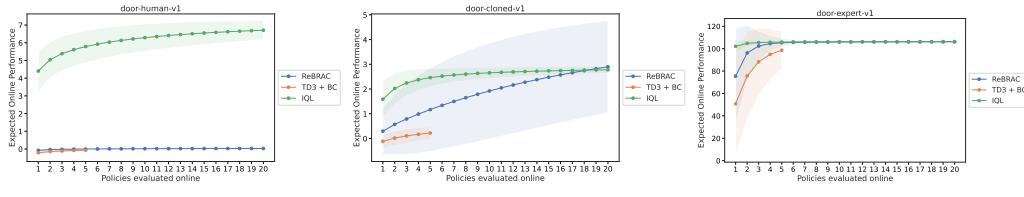


Figure 8: TD3+BC, IQL and ReBRAC visualised Expected Online Performance under uniform policy selection on Door tasks.

Table 28: TD3+BC, IQL and ReBRAC Expected Online Performance under uniform policy selection on Hammer tasks.

Policies	human			cloned			expert		
	TD3+BC	IQL	ReBRAC	TD3+BC	IQL	ReBRAC	TD3+BC	IQL	ReBRAC
1	0.2 ± 0.0	1.6 ± 0.4	0.3 ± 0.2	1.1 ± 0.8	1.9 ± 1.1	4.5 ± 5.5	60.2 ± 55.4	128.7 ± 1.1	101.9 ± 49.9
2	0.3 ± 0.0	1.8 ± 0.4	0.3 ± 0.2	1.6 ± 0.8	2.5 ± 1.1	7.3 ± 6.1	90.0 ± 48.4	129.3 ± 0.9	124.3 ± 26.8
3	0.3 ± 0.0	1.9 ± 0.4	0.4 ± 0.2	1.8 ± 0.7	2.8 ± 1.0	9.2 ± 6.2	105.7 ± 38.1	129.6 ± 0.7	130.1 ± 13.6
4	0.3 ± 0.0	2.0 ± 0.4	0.4 ± 0.2	2.0 ± 0.6	3.1 ± 0.9	10.6 ± 6.2	114.3 ± 29.3	129.7 ± 0.5	131.9 ± 7.0
5	0.3 ± 0.0	2.1 ± 0.4	0.5 ± 0.2	2.1 ± 0.6	3.3 ± 0.9	11.7 ± 6.2	119.2 ± 22.4	129.8 ± 0.4	132.6 ± 3.8
6	-	2.1 ± 0.4	0.5 ± 0.2	-	3.4 ± 0.8	12.6 ± 6.1	-	129.9 ± 0.3	133.0 ± 2.4
7	-	2.2 ± 0.4	0.5 ± 0.2	-	3.5 ± 0.7	13.3 ± 6.0	-	129.9 ± 0.2	133.2 ± 1.7
8	-	2.2 ± 0.4	0.5 ± 0.2	-	3.6 ± 0.7	14.0 ± 5.9	-	129.9 ± 0.2	133.4 ± 1.4
9	-	2.3 ± 0.3	0.6 ± 0.2	-	3.7 ± 0.6	14.6 ± 5.8	-	130.0 ± 0.1	133.5 ± 1.2
10	-	2.3 ± 0.3	0.6 ± 0.2	-	3.7 ± 0.6	15.1 ± 5.7	-	130.0 ± 0.1	133.6 ± 1.1
11	-	2.3 ± 0.3	0.6 ± 0.1	-	3.8 ± 0.5	15.5 ± 5.6	-	130.0 ± 0.1	133.7 ± 1.0
12	-	2.3 ± 0.3	0.6 ± 0.1	-	3.8 ± 0.5	16.0 ± 5.4	-	130.0 ± 0.1	133.8 ± 0.9
13	-	2.4 ± 0.3	0.6 ± 0.1	-	3.9 ± 0.5	16.3 ± 5.3	-	130.0 ± 0.1	133.9 ± 0.8
14	-	2.4 ± 0.3	0.6 ± 0.1	-	3.9 ± 0.4	16.7 ± 5.2	-	130.0 ± 0.1	133.9 ± 0.7
15	-	2.4 ± 0.3	0.6 ± 0.1	-	3.9 ± 0.4	17.0 ± 5.1	-	130.0 ± 0.1	134.0 ± 0.6
16	-	2.4 ± 0.3	0.6 ± 0.1	-	4.0 ± 0.4	17.3 ± 4.9	-	130.0 ± 0.1	134.0 ± 0.6
17	-	2.4 ± 0.2	0.6 ± 0.1	-	4.0 ± 0.4	17.5 ± 4.8	-	130.0 ± 0.0	134.0 ± 0.5
18	-	2.4 ± 0.2	0.6 ± 0.1	-	4.0 ± 0.3	17.8 ± 4.7	-	130.0 ± 0.0	134.0 ± 0.5
19	-	2.5 ± 0.2	0.6 ± 0.1	-	4.0 ± 0.3	18.0 ± 4.6	-	130.0 ± 0.0	134.1 ± 0.4
20	-	2.5 ± 0.2	0.7 ± 0.1	-	4.0 ± 0.3	18.2 ± 4.5	-	130.0 ± 0.0	134.1 ± 0.4

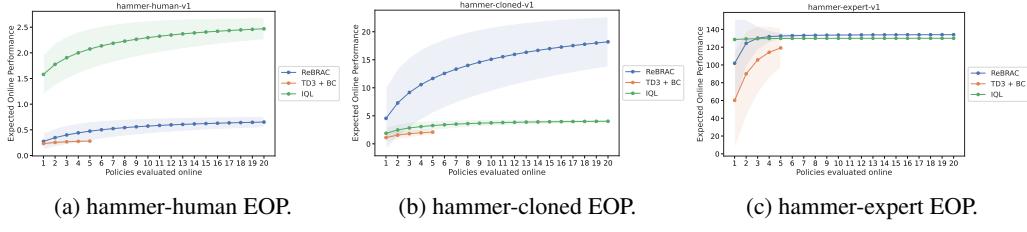


Figure 9: TD3+BC, IQL and ReBRAC visualised Expected Online Performance under uniform policy selection on Hammer tasks.

Table 29: TD3+BC, IQL and ReBRAC Expected Online Performance under uniform policy selection on tasks.

Policies	human			cloned			expert		
	TD3+BC	IQL	ReBRAC	TD3+BC	IQL	ReBRAC	TD3+BC	IQL	ReBRAC
1	-0.2 ± 0.1	0.2 ± 0.2	-0.1 ± 0.1	-0.2 ± 0.1	-0.0 ± 0.1	0.5 ± 0.8	21.4 ± 43.2	106.0 ± 1.4	73.5 ± 44.3
2	-0.2 ± 0.1	0.2 ± 0.2	-0.1 ± 0.1	-0.2 ± 0.1	0.0 ± 0.1	0.9 ± 0.9	38.8 ± 51.9	106.8 ± 1.0	96.0 ± 27.4
3	-0.1 ± 0.1	0.3 ± 0.2	-0.0 ± 0.0	-0.2 ± 0.1	0.1 ± 0.1	1.2 ± 0.9	52.6 ± 54.0	107.2 ± 0.8	103.7 ± 15.9
4	-0.1 ± 0.1	0.3 ± 0.2	-0.0 ± 0.0	-0.1 ± 0.1	0.1 ± 0.1	1.4 ± 0.9	63.7 ± 53.1	107.4 ± 0.7	106.7 ± 9.4
5	-0.1 ± 0.1	0.4 ± 0.2	-0.0 ± 0.0	-0.1 ± 0.0	0.1 ± 0.1	1.6 ± 0.9	72.5 ± 50.7	107.5 ± 0.6	107.9 ± 5.8
6	-	0.4 ± 0.2	-0.0 ± 0.0	-	0.1 ± 0.1	1.7 ± 0.8	-	107.6 ± 0.5	108.6 ± 3.8
7	-	0.4 ± 0.2	-0.0 ± 0.0	-	0.1 ± 0.1	1.8 ± 0.8	-	107.7 ± 0.5	108.9 ± 2.6
8	-	0.4 ± 0.2	0.0 ± 0.0	-	0.1 ± 0.1	1.9 ± 0.8	-	107.7 ± 0.5	109.2 ± 1.9
9	-	0.5 ± 0.2	0.0 ± 0.0	-	0.1 ± 0.1	2.0 ± 0.7	-	107.8 ± 0.4	109.3 ± 1.5
10	-	0.5 ± 0.2	0.0 ± 0.0	-	0.1 ± 0.1	2.1 ± 0.7	-	107.8 ± 0.4	109.4 ± 1.3
11	-	0.5 ± 0.2	0.0 ± 0.0	-	0.2 ± 0.1	2.1 ± 0.7	-	107.8 ± 0.4	109.5 ± 1.1
12	-	0.5 ± 0.2	0.0 ± 0.0	-	0.2 ± 0.1	2.2 ± 0.6	-	107.9 ± 0.4	109.6 ± 1.0
13	-	0.5 ± 0.2	0.0 ± 0.0	-	0.2 ± 0.1	2.2 ± 0.6	-	107.9 ± 0.4	109.7 ± 0.9
14	-	0.5 ± 0.2	0.0 ± 0.0	-	0.2 ± 0.1	2.3 ± 0.6	-	107.9 ± 0.4	109.8 ± 0.9
15	-	0.5 ± 0.2	0.0 ± 0.0	-	0.2 ± 0.1	2.3 ± 0.6	-	107.9 ± 0.4	109.8 ± 0.8
16	-	0.5 ± 0.2	0.0 ± 0.0	-	0.2 ± 0.1	2.3 ± 0.5	-	108.0 ± 0.4	109.9 ± 0.8
17	-	0.6 ± 0.2	0.0 ± 0.0	-	0.2 ± 0.1	2.4 ± 0.5	-	108.0 ± 0.4	109.9 ± 0.8
18	-	0.6 ± 0.2	0.0 ± 0.0	-	0.2 ± 0.1	2.4 ± 0.5	-	108.0 ± 0.3	109.9 ± 0.8
19	-	0.6 ± 0.2	0.0 ± 0.0	-	0.2 ± 0.1	2.4 ± 0.5	-	108.0 ± 0.3	110.0 ± 0.7
20	-	0.6 ± 0.1	0.0 ± 0.0	-	0.2 ± 0.1	2.5 ± 0.4	-	108.0 ± 0.3	110.0 ± 0.7

Table 34: ReBRAC ablations for pen tasks. We report final normalized score averaged over 4 unseen training seeds.

Ablation	human	cloned	expert	Average
TD3+BC, paper	0.0	0.0	0.3	0.0
TD3+BC, our	65.9 ± 24.6	78.1 ± 5.7	144.9 ± 7.5	96.3
TD3+BC, tuned	77.6 ± 18.5 (-23.9%)	78.1 ± 5.7 (-8.5%)	144.9 ± 7.5 (-10.3%)	100.2 (-12.5%)
ReBRAC w/o LN	78.6 ± 14.8 (-22.9%)	21.3 ± 11.0 (-75.1%)	86.7 ± 59.8 (-44.6%)	62.1 (-45.8%)
ReBRAC w/o layer	89.1 ± 14.7 (-12.6%)	106.7 ± 13.9 (+24.9%)	147.2 ± 5.7 (-6.0%)	114.3 (-0.3%)
ReBRAC w/o actor penalty	-0.5 ± 1.3 (-100.5%)	0.6 ± 1.6 (-99.3%)	0.0 ± 3.6 (-100.0%)	0.0 (-100.0%)
ReBRAC w/o critic penalty	99.9 ± 6.1 (-2.1%)	75.0 ± 16.7 (-12.2%)	154.6 ± 1.8 (+1.3%)	109.8 (4.2%)
ReBRAC w large batch	67.2 ± 9.0 (-34.1%)	83.2 ± 23.3 (-2.6%)	155.0 ± 6.8 (-1.0%)	101.8 (-11.2%)
ReBRAC	102.0 ± 10.8	85.4 ± 24.2	156.6 ± 1.4	114.6

Table 35: ReBRAC ablations for door tasks. We report final normalized score averaged over 4 unseen training seeds.

Ablation	human	cloned	expert	Average
TD3+BC, paper	0.0	0.0	0.0	0.0
TD3+BC, our	0.0 ± 0.1	0.4 ± 1.0	102.5 ± 2.9	34.3
TD3+BC, tuned	0.0 ± 0.1 (-)	0.4 ± 1.0 (+100.0%)	105.8 ± 0.3 (+0.8%)	35.4 (+1.1%)
ReBRAC w/o LN	-0.1 ± 0.0 (-)	-0.3 ± 0.0 (-250.0%)	106.0 ± 0.8 (+1.0%)	35.1 (+0.3%)
ReBRAC w/o layer	0.0 ± 0.0 (-)	0.1 ± 0.5 (-50.0%)	104.4 ± 2.3 (-0.5%)	34.8 (-0.6%)
ReBRAC w/o actor penalty	-0.1 ± 0.1 (-)	0.0 ± 0.0 (-100.0%)	0.0 ± 0.2 (-100.0%)	0.0 (-100.0%)
ReBRAC w/o critic penalty	0.0 ± 0.0 (-)	0.1 ± 0.0 (-50.0%)	106.1 ± 0.3 (+1.1%)	35.4 (+1.1%)
ReBRAC w large batch	-0.1 ± 0.1 (-)	0.1 ± 0.3 (-50.0%)	106.1 ± 0.1 (+1.1%)	35.3 (+0.9%)
ReBRAC	0.0 ± 0.0	0.2 ± 0.3	104.9 ± 2.2	35.0

Table 36: ReBRAC ablations for hammer tasks. We report final normalized score averaged over 4 unseen training seeds.

Ablation	human	cloned	expert	Average
TD3+BC, paper	0.0	0.0	0.0	0.0
TD3+BC, our	0.3 ± 0.4	1.1 ± 1.1	127.0 ± 0.4	42.8
TD3+BC, tuned	0.3 ± 0.4 (+50.0%)	1.1 ± 1.1 (-80.0%)	127.0 ± 0.4 (-5.3%)	42.8 (-8.1%)
ReBRAC w/o LN	0.2 ± 0.0 (0.0%)	1.0 ± 1.0 (-81.8%)	9.9 ± 19.1 (-92.6%)	3.6 (-92.3%)
ReBRAC w/o layer	0.1 ± 0.0 (-50.0%)	21.3 ± 19.7 (+287.3%)	133.1 ± 0.5 (-0.8%)	51.5 (+10.5%)
ReBRAC w/o actor penalty	0.0 ± 0.0 (-100.0%)	0.0 ± 0.1 (-100.0%)	0.0 ± 0.1 (-100.0%)	0.0 (-100.0%)
ReBRAC w/o critic penalty	0.1 ± 0.1 (-50.0%)	1.9 ± 0.7 (-65.5%)	134.1 ± 0.2 (-0.1%)	45.3 (-2.8%)
ReBRAC w large batch	0.3 ± 0.8 (+50.0%)	10.6 ± 14.0 (+92.7%)	133.4 ± 0.5 (-0.6%)	48.1 (+3.2%)
ReBRAC	0.2 ± 0.2	5.5 ± 2.5	134.2 ± 0.4	46.6

Table 37: ReBRAC ablations for relocate tasks. We report final normalized score averaged over 4 unseen training seeds.

Ablation	human	cloned	expert	Average
TD3+BC, paper	0.0	0.0	0.0	0.0
TD3+BC, our	0.0 ± 0.0	-0.1 ± 0.0	107.9 ± 0.6	35.9
TD3+BC, tuned	0.0 ± 0.0 (-)	-0.1 ± 0.0 (-105.3%)	107.9 ± 0.6 (+1.2%)	35.9 (-0.5%)
ReBRAC w/o LN	-0.2 ± 0.0 (-)	0.0 ± 0.3 (-100.0%)	-0.1 ± 0.0 (-100.1%)	-0.1 (-100.3%)
ReBRAC w/o layer	0.1 ± 0.3 (-)	1.7 ± 2.1 (-10.5%)	105.0 ± 3.1 (-1.5%)	35.6 (-1.4%)
ReBRAC w/o actor penalty	-0.1 ± 0.0 (-)	0.0 ± 0.0 (-100.0%)	-0.1 ± 0.1 (-100.1%)	0.0 (-100.0%)
ReBRAC w/o critic penalty	0.0 ± 0.1 (-)	1.9 ± 1.9 (0.0%)	109.6 ± 1.2 (+2.8%)	37.1 (+2.8%)
ReBRAC w large batch	0.0 ± 0.0 (-)	0.1 ± 0.2 (-94.7%)	109.6 ± 0.9 (+2.8%)	36.5 (+1.1%)
ReBRAC	0.0 ± 0.0	1.9 ± 2.3	106.6 ± 3.1	36.1