Figure 6: Revisiting Fig. 1 for the discounted cases where $\gamma \in (0, 1)$.

430 A MORE EXAMPLES WITH MULTIPLE FIXED POINTS

431 A.1 CASES OF $\gamma = 1$

432 Consider MDP examples with an terminal state 0, as shown in Fig. 1 (we adapt dynamic programming
433 examples [5, 39] into reinforcement learning settings.),

- 434 • **Shortest path problem (deterministic)** in Fig. 1(a): At state 1, an agent transits to either state
435 1 or 0 with reward 0 or b , respectively. Assume the value function for state 0 is $V(0) = 0$. The
436 Bellman's optimality equation for state 1 is $V(1) = \max\{V(1), b\}$, where any $V(1) \geq b$ is a
437 feasible solution. If initialize $V_0(1) \geq b$, a resulting policy is that an agent at state 1 always transits
438 back to state 1; otherwise, drives to terminal state 0 (always returns back to itself with reward 0).
- 439 • **Blackmailer's problem (stochastic)** in Fig. 1(b): At state 1, a profit maximizing blackmailer
440 demands a cash amount $a \in (0, 1]$; a victim transits to state 1 with probability a or state 0 with
441 probability $1 - a$, respectively. At state 0, a victim always refuses to yield, i.e., $V(0) = 0$. The
442 Bellman's optimality equation for state 1 is $V(1) = \max_a\{a + (1 - a)V(1)\}$, where any $V(1) \geq 1$
443 is a feasible solution. If initialize $V_0(1) > 1$, the blackmailer's policy is demanding $a \rightarrow 0$ to keep
444 the victim at state 1; otherwise, demanding $a = 1$ that drives the victim to terminal state 0.
- 445 • **Optimal stopping problem (terminating policies)** in Fig. 1(c): In a space \mathbb{R}^2 with terminal
446 state of point 0, an agent at point $x \neq 0$ moves to either point 0 with negative reward $-c$ or
447 point αx with reward $-||x||$, respectively, where $\alpha \in (0, 1)$. The Bellman's optimality equation
448 is $V(x) = \max\{-c, -||x|| + V(\alpha x)\}$ and the optimal policy is to continue inside the sphere of
449 radius $(1 - \alpha)c$ and to stop outside. If add a cone region C within which an agent always receives
450 a reward $-c$, a second policy is jumping to point 0 at any point in region C .

451 A.2 CASES OF $\gamma \in (0, 1)$

452 First, we consider the discounted formulations of the three examples (shown in Fig. 1), as shown in
453 Fig. 6 where $\gamma \in (0, 1)$. The differences are marked in **red**.

- 454 • (a) **Shortest path problem (deterministic, discounted case)**: Given two states 1 and 0, an agent
455 at state 1 transits to either state 1 or 0 with rewards $r = c$ or $r = b$, respectively. $c > (1 - \gamma) \cdot b$.
456 At state 0, the value function is $V(0) = 0$. At state 1, the Bellman's optimality equation is
457 $V(1) = \max\{c + \gamma \cdot V(1), b\}$, where any $V(1) \geq (b - c)/\gamma$ is a solution. If initialize $V_0(1) \geq b$,
458 an agent obtains a policy that always transits back to state 1; otherwise, a result policy drives to
459 terminal state 0.
- 460 • (b) **Blackmailer's problem (stochastic, discounted case)**: Different from (a), a profit maximizing
461 blackmailer/agent at 1 demands a cash amount $a \in (0, 1]$ (an action), while a victim transits to state
462 1 with probability a or to state 0 with probability $1 - a$, respectively. At state 0, a victim always
463 refuses to yield to the blackmailer's demand, i.e., $V(0) = 0$. The Bellman's optimality equation
464 is $V(1) = \max_a\{a + \gamma \cdot (1 - a)V(1)\}$ for state 1, where any $V(1) \geq 1$ is a feasible solution. If
465 initialize $V_0(1) = c > 1$, the blackmailer's policy is demanding $a \rightarrow 0$ at the k -th step to keep
466 the victim stay at state 1, for any $k \leq K_0 = -\lfloor \log_\gamma c \rfloor$; and taking $a = 1$ to transit to terminal

467 state 0 at the k -th step, for any $k \geq K_0 + 1$; otherwise initialize $V_0(1) = c \leq 1$, the result policy is
 468 demanding the maximum $a = 1$ that drives the victim to a refusal state 0 (a terminal state).

469 • (c) **Optimal stopping problem (terminating policies, discounted case)**: In a space \mathbb{R}^2 with
 470 terminating state at point 0, at point $x \neq 0$ an agent moves to either point 0 with negative reward $-c$
 471 or point αx with reward $-||x||$, respectively, where $\alpha \in (0, 1)$. The Bellman's optimality equation
 472 is $V(x) = \max\{-c, -||x|| + \gamma \cdot V(\alpha x)\}$ and the optimal policy is to continue inside the sphere of
 473 radius $(1 - \alpha)c$ and to stop outside. If add a cone region C within which an agent always receives
 474 a reward $-c$, a second policy is jumping to point 0 at any point in region C .

475 Then, we elaborate how the proposed H-term fixes the problems in Fig. 6.

476 (a) **Shortest path problem (deterministic, discounted case)**

477 Assume $V_0(1) \geq b$ and $c > (1 - \gamma)b$, we have

$$\begin{aligned}
 V_1(1) &= c + \gamma \cdot V_0(1) \geq c + \gamma \cdot b > b \\
 V_2(1) &= c + \gamma \cdot c + \gamma^2 \cdot V_0(1) \geq (1 + \gamma)c + \gamma^2 b > b \\
 V_3(1) &= c + \gamma \cdot c + \gamma^2 c + \gamma^3 \cdot V_0(1) \geq (1 + \gamma + \gamma^2)c + \gamma^3 b > b \\
 &\dots \\
 V_k(1) &= \sum_{i=0}^{k-1} \gamma^i \cdot c + \gamma^k \cdot V_0(1) \geq \sum_{i=0}^{k-1} \gamma^i \cdot c + \gamma^k b > b \\
 &\dots \\
 V^*(1) &= \sum_{i=0}^{\infty} \gamma^i \cdot c = \frac{1}{1 - \gamma} c > b
 \end{aligned} \tag{9}$$

478 The values of $H(0)$ and $H(1)$ are as follows:

$$H(0) = 0, \quad H(1) = -b - \sum_{k=2}^{\infty} \left(\sum_{i=1}^{k-1} \gamma^{i-1} \cdot c + \gamma^k b \right) = -\infty. \tag{10}$$

479 Adding the above H-values to state 1 and 0, respectively, we have

$$\begin{aligned}
 V^*(1) + H(1) &= \sum_{i=0}^{\infty} \gamma^i \cdot c - \infty = -\infty \\
 V^*(0) + H(0) &= b.
 \end{aligned} \tag{11}$$

480 Therefore, $V^*(1) + H(1) < V^*(0) + H(0)$, independent of the initial value $V_0(1)$. That is, an agent
 481 always obtains a policy that drives to terminal state 0 at step 1.

482 (b) **Blackmailer's problem (stochastic, discounted case)**

483 If initialize $V_0(1) = c > 1$, the blackmailer's policy is demanding $a \rightarrow 0$ at the k -th step to keep the
 484 victim stay at state 1, for any $k \leq K_0 = -\lfloor \log_{\gamma} c \rfloor$; and taking $a = 1$ to transit to terminal state 0 at
 485 the k -th step, for any $k \geq K_0 + 1$; otherwise initialize $V_0(1) = c \leq 1$, the result policy is demanding
 486 the maximum $a = 1$ that drives the victim to a refusal state 0 (a terminal state).

487 The values of $H(0)$ and $H(1)$ are as follows:

$$H(0) = 0, \quad H(1) = - \sum_{k=1}^{\infty} \sum_{i=1}^{k-1} \gamma^{i-1} \cdot a = -\infty. \tag{12}$$

488 For arbitrary initial value of $V_0(1)$, $V_1(1) = a + (1 - a) \cdot \gamma(V_0(1) + H(1))$ take maximum $V_1(1) = 1$
 489 when $a = 1$. Therefore, the policy always drives to terminal state 0 at step 1.

490 (c) **Optimal stopping problem (terminating policies, discounted case)**

491 Any policy that takes infinite steps will have

$$H(x) = -c - \sum_{k=2}^{\infty} \left[\sum_{i=1}^{k-1} \gamma^i \cdot \alpha^i \cdot ||x|| + \gamma^k \cdot (-c) \right] = -\infty \tag{13}$$

492 and a direct jumping policy will have $H(x) = -c$. Therefore, the H-term drives to a terminating
493 policy.

494 **B MUJoCo TASKS WITH MULTIPLE POLICIES**495 **B.1 DESCRIPTION OF MUJoCo TASKS**

496 We selected six challenging robotic locomotion tasks from MuJoCo, namely, Swimmer-v3, Hopper-
 497 v3, Walker2D-v3, HalfCheetah-v3, Ant-v3, Humanoid-v3. Table 3 lists the action space and state
 498 space of each task.

Table 3: The state and action spaces of six challenging MuJoCo tasks.

Tasks	Agent	Action Space	State Space
Swimmer-v3	Three-link swimming robot	2	8
Hopper-v3	Two-dimensional one-legged robot	3	11
Walker2d-v3	Two-dimensional bipedal robot	6	17
HalfCheetah-v3	Two-dimensional robot	6	17
Ant-v3	Four-legged creature	8	111
Humanoid-v3	Three-dimensional bipedal robot	17	376

499 **B.2 MULTIPLE POLICIES IN MUJoCo TASKS**

500 In the supplementary files, we includes rendered videos of different policies, as given in Table 4.

- 501 • Different policies are obtained over 20 runs of the PPO algorithm. We rendered theses polices and
 502 classified them by physical gaits.
- 503 • The policies in bold texts are physically stationary.

Table 4: List of video files for different policies.

Task	Different Policies	Video Name
Hopper	hopping	hopper_hopping.mp4
	diving	hopper_diving.mp4
	standing	hopper_standing.mp4
Ant	running	ant_running.mp4
	standing	ant_standing.mp4
	flipping	ant_flipping.mp4
Walker	walking	walker_walking.mp4
	diving	walker_diving.mp4
	standing	walker_standing.mp4
Humanoid	two-legs	humanoid_two_legs.mp4
	one-leg	humanoid_one_leg.mp4
	backward	humanoid_backward.mp4
HalfCheetah	running	halfcheetah_running.mp4
	diving	halfcheetah_diving.mp4
	flipping	halfcheetah_flipping.mp4
	standing	halfcheetah_standing.mp4
Swimmer	moving	swimmer_moving.mp4
	standing	swimmer_standing.mp4

504 C REINFORCEMENT LEARNING AND BELLMAN EQUATION

505 A reinforcement learning (RL) [44] agent interacts with an unknown environment and learns an
 506 optimal policy that maximizes the cumulative reward. Mathematically, the environment can be
 507 formulated as a Markov Decision Process (MDP) with the five-tuple $\langle \mathcal{S}, \mathcal{A}, \mathbb{P}, R, \gamma \rangle$. Here \mathcal{S} and
 508 \mathcal{A} denote the state and action spaces; $\mathbb{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ denotes a transition probability function,
 509 where Δ is a probability simplex; $R : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ denotes a reward function; and $\gamma \in (0, 1]$
 510 denotes a discount factor. The objective is to find an optimal policy $\pi^* : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ that maximizes
 511 (discounted) expected reward.

512 Consider a discrete, finite, discounted MDP with infinite horizon, one can define the Q-value function
 513 of a state-action pair (s, a) under policy π as follows

$$Q^\pi(s, a) = \mathbb{E}_{S_{k+1} \sim \mathbb{P}(\cdot | S_k, A_k), A_{k+1} \sim \pi(S_{k+1}, \cdot)} \left[\sum_{k=0}^{\infty} \gamma^k \cdot R(S_k, A_k, S_{k+1}) | S_0 = s, A_0 = a \right], \quad (14)$$

514 where $R(S_k, A_k, S_{k+1})$ denotes the immediate reward when taking action A_k at state S_k and arriving
 515 at state S_{k+1} , capital letters denote random variables and lowercase letters denote values. The
 516 conventional objective function $J(\theta)$ of reinforcement learning [44] takes the following form

$$J(\theta) \triangleq \mathbb{E}_{S_0, A_0} [Q^{\pi_\theta}(S_0, A_0)] = \mathbb{E}_{\tau \sim \pi} [R(\tau) \cdot P(\tau | \pi_\theta)], \quad (15)$$

517 where τ is a trajectory, i.e., $\tau = (S_0, A_0, \dots)$, and

$$518 P(\tau | \pi_\theta) = d_0(s_0) \cdot \prod_{k=0}^T \mathbb{P}(s_{k+1} | s_k, a_k) \pi_\theta(a_k | s_k).$$

519 The Bellman equation [44] converts (14) into a recursive form as follows

$$\begin{aligned} Q^\pi(s, a) &= \sum_{s' \in \mathcal{S}} \mathbb{P}(s' | s, a) \left[R(s, a, s') + \gamma \sum_{a' \in \mathcal{A}} \pi(s', a') Q^\pi(s', a') \right] \\ &= R(s, a) + \gamma \sum_{s' \in \mathcal{S}} \mathbb{P}(s' | s, a) \sum_{a' \in \mathcal{A}} \pi(s', a') Q^\pi(s', a'), \end{aligned} \quad (16)$$

520 which expresses the expected reward as a summation of immediate reward $R(s, a)$ and discounted fu-
 521 ture rewards, and the immediate reward $R(s, a)$ is defined as $R(s, a) = \sum_{s' \in \mathcal{S}} \mathbb{P}(s' | s, a) R(s, a, s')$.

522 The Bellman's optimality equation [44] is

$$Q^*(s, a) = \sum_{s' \in \mathcal{S}} \mathbb{P}(s' | s, a) \left[R(s, a, s') + \gamma \max_{a'} Q^*(s', a') \right]. \quad (17)$$

523 The optimal policy π^* is given by a greedy strategy such that $\pi^* = \arg \max_{\pi} Q^\pi(s, a)$.

524 **D QUANTUM K-SPIN HAMILTONIAN FORMULATION OF REINFORCEMENT**
 525 **LEARNING**

526 We provide the detailed steps of reformulating (14) into a K -spin Hamiltonian equation

$$\begin{aligned}
 H(\theta) &\triangleq -\mathbb{E}_{S_0, A_0} [Q^{\pi_\theta}(S_0, A_0)] \\
 &= -\mathbb{E}_{S_0, A_k \sim \pi_\theta(S_k, \cdot), S_{k+1} \sim \mathbb{P}(\cdot | S_k, A_k)} \left[\sum_{k=0}^{\infty} \gamma^k \cdot R(S_k, A_k) \right] \\
 &= -\sum_{k=0}^{K-1} \mathbb{E}_{S_0, A_0, \dots, S_k \sim \mathbb{P}(\cdot | S_{k-1}, A_{k-1}), A_k \sim \pi_\theta(S_k, \cdot)} [\gamma^k \cdot R(S_k, A_k)] \\
 &= -\sum_{k=0}^{K-1} \sum_{\mu_0}^{S \times A} \dots \sum_{\mu_k}^{S \times A} \gamma^k \cdot R(\mu_k) \cdot d_0(S_0) \cdot \pi_\theta(\mu_0) \prod_{i=0}^{k-1} [\mathbb{P}(S_{i+1} | \mu_i) \cdot \pi_\theta(\mu_{i+1})] \quad (18) \\
 &= -\sum_{k=0}^{K-1} \sum_{\mu_0}^{S \times A} \dots \sum_{\mu_k}^{S \times A} \left[\gamma^k \cdot R(\mu_k) \cdot d_0(S_0) \cdot \prod_{i=0}^{k-1} \mathbb{P}(S_{i+1} | \mu_i) \right] \cdot \pi_\theta(\mu_0) \dots \pi_\theta(\mu_k) \\
 &= -\sum_{k=0}^{K-1} \sum_{\mu_0}^{S \times A} \dots \sum_{\mu_k}^{S \times A} L_{\mu_0, \dots, \mu_k} \cdot \pi_\theta(\mu_0) \dots \pi_\theta(\mu_k),
 \end{aligned}$$

527 where $K \rightarrow \infty$, and the density function is

$$L_{\mu_0, \dots, \mu_k} = \gamma^k \cdot R(\mu_k) \cdot d_0(S_0) \cdot \prod_{i=0}^{k-1} \mathbb{P}(S_{i+1} | \mu_i). \quad (19)$$

528 E DERIVATION STEPS FOR HAMILTONIAN’S POLICY GRADIENTS

529 We provide the policy gradient of the quantum K-spin Hamiltonian equation in (3) for both stochastic
530 and deterministic cases, which are variants of the well-known policy gradient theorem [44].

531 **Theorem 3. (Hamiltonian’s stochastic policy gradient)** *The stochastic gradient of the K-spin*
532 *Hamiltonian equation (3) w.r.t. parameter θ is*

$$\nabla_{\theta} H(\theta) = -\mathbb{E}_{\mu_0, \dots, \mu_{K-1}} \left[\sum_{k=0}^{K-1} \gamma^k \cdot R(\mu_k) \cdot \nabla_{\theta} \log (\pi_{\theta}(\mu_0) \cdot \pi_{\theta}(\mu_1) \cdots \pi_{\theta}(\mu_k)) \right]. \quad (20)$$

533 **Corollary 1.** *When $K \rightarrow \infty$, the Hamiltonian’s stochastic policy gradient $\nabla_{\theta} H(\theta)$ in (20) is equal*
534 *to the stochastic policy gradient $\nabla_{\theta} J(\theta)$ in [45],*

$$\lim_{K \rightarrow \infty} \nabla_{\theta} H(\theta) = -\nabla_{\theta} J(\theta) = -\mathbb{E}_{s \sim d_{\theta}, a \sim \pi_{\theta}} [Q^{\pi_{\theta}}(s, a) \nabla_{\theta} \log \pi_{\theta}(s, a)]. \quad (21)$$

535 Let $\eta_{\theta}(\cdot) : \mathcal{S} \rightarrow \mathcal{A}$ denote a deterministic policy, while we use $\tilde{\pi}_{\theta, \delta}(\mu)$ to represent that a Gaussian
536 noise (a.k.a, an exploration noise) with standard deviation $\delta > 0$ is added in the exploration process.

537 **Theorem 4. (Hamiltonian’s deterministic policy gradient)** *The deterministic gradient of the K-spin*
538 *Hamiltonian equation (3) w.r.t. parameter θ is*

$$\nabla_{\theta} H'(\theta) = -\mathbb{E}_{\mu_0, \dots, \mu_{K-1}} \left[\sum_{k=0}^{K-1} \gamma^k \cdot R(\mu_k) \cdot \nabla_{\theta} \log (\tilde{\pi}_{\theta, \delta}(\mu_0) \cdot \tilde{\pi}_{\theta, \delta}(\mu_1) \cdots \tilde{\pi}_{\theta, \delta}(\mu_k)) \right]. \quad (22)$$

539 **Corollary 2.** *When $K \rightarrow \infty$, the Hamiltonian’s deterministic policy gradient $\nabla_{\theta} H'(\theta)$ in (22) is*
540 *equal to the deterministic policy gradient $\nabla_{\theta} J'(\theta)$ in [43],*

$$\lim_{K \rightarrow \infty} \nabla_{\theta} H'(\theta) = -\nabla_{\theta} J'(\theta) = -\mathbb{E}_{s \sim d_{\theta}} \left[\nabla_a Q^{\tilde{\pi}_{\theta, \delta}}(s, a) \Big|_{a=\eta_{\theta}} \nabla_{\theta} \eta_{\theta}(s) \right]. \quad (23)$$

541 **Corollary 3.** *When the variance of the exploration noise approaches zero, i.e., $\delta \rightarrow 0$, the deter-*
542 *ministic policy gradient $\nabla_{\theta} H'(\theta)$ is the limiting case of the stochastic policy gradient $\nabla_{\theta} H(\theta)$,*
543

$$\nabla_{\theta} H'(\theta) = \lim_{\delta \rightarrow 0} \nabla_{\theta} H(\theta). \quad (24)$$

544 E.1 PROOF OF THEOREM 3: HAMILTONIAN’S STOCHASTIC POLICY GRADIENT

Proof.

$$\begin{aligned} \nabla_{\theta} H(\theta) &= - \sum_{k=0}^{K-1} \sum_{\mu_0}^{\mathcal{S} \times \mathcal{A}} \cdots \sum_{\mu_k}^{\mathcal{S} \times \mathcal{A}} L_{\mu_0, \dots, \mu_k} \nabla_{\theta} [\pi_{\theta}(\mu_0) \cdots \pi_{\theta}(\mu_k)] \\ &= - \sum_{k=0}^{K-1} \sum_{\mu_0}^{\mathcal{S} \times \mathcal{A}} \cdots \sum_{\mu_k}^{\mathcal{S} \times \mathcal{A}} L_{\mu_0, \dots, \mu_k} [\pi_{\theta}(\mu_0) \cdots \pi_{\theta}(\mu_k)] \nabla_{\theta} \log [\pi_{\theta}(\mu_0) \cdots \pi_{\theta}(\mu_k)] \\ &= - \sum_{k=0}^{K-1} \sum_{\mu_0}^{\mathcal{S} \times \mathcal{A}} \cdots \sum_{\mu_k}^{\mathcal{S} \times \mathcal{A}} \gamma^k \cdot R(\mu_k) \cdot d_0(S_0) \cdot \pi_{\theta}(\mu_0) \prod_{i=0}^{k-1} [\mathbb{P}(S_{i+1} | \mu_i) \cdot \pi_{\theta}(\mu_{i+1})] \cdot \nabla_{\theta} \log [\pi_{\theta}(\mu_0) \cdots \pi_{\theta}(\mu_k)] \\ &= -\mathbb{E}_{\mu_0, \dots, \mu_{K-1}} \left[\sum_{k=0}^{K-1} \gamma^k \cdot R(\mu_k) \cdot \nabla_{\theta} \log [\pi_{\theta}(\mu_0) \cdots \pi_{\theta}(\mu_k)] \right], \end{aligned} \quad (25)$$

545 where $\mu_k = (S_k, A_k)$, $S_0 \sim d_0(\cdot)$, $A_k \sim \pi_{\theta}(S_k, \cdot)$, $S_{k+1} \sim \mathbb{P}(\cdot | S_k, A_k)$ for $k = 0 \cdots K$. \square

546 E.2 PROOF OF COROLLARY 1

Proof.

$$\begin{aligned}
\nabla_{\theta} H(\theta) &\stackrel{(a)}{=} - \sum_{k=0}^{K-1} \sum_{\mu_0}^{S \times A} \cdots \sum_{\mu_k}^{S \times A} L_{\mu_0, \dots, \mu_k} \nabla_{\theta} [\pi_{\theta}(\mu_0) \cdots \pi_{\theta}(\mu_k)] \\
&\stackrel{(b)}{=} - \sum_{k=0}^{K-1} \sum_{\mu_0}^{S \times A} \cdots \sum_{\mu_k}^{S \times A} L_{\mu_0, \dots, \mu_k} \sum_{i=0}^k \pi_{\theta}(\mu_0) \cdots \pi_{\theta}(\mu_{i-1}) \pi_{\theta}(\mu_{i+1}) \cdots \pi_{\theta}(\mu_k) \nabla_{\theta} \pi_{\theta}(\mu_i) \\
&\stackrel{(c)}{=} - \sum_{k=0}^{K-1} \sum_{\mu_0}^{S \times A} \cdots \sum_{\mu_k}^{S \times A} \gamma^k \cdot R(\mu_k) \cdot d_0(S_0) \prod_{i=0}^{k-1} \mathbb{P}(S_{i+1} | \mu_i) \sum_{i=0}^k \left[\prod_{j=0}^{i-1} \pi_{\theta}(\mu_j) \cdot \nabla_{\theta} \pi_{\theta}(\mu_i) \cdot \prod_{j=i+1}^k \pi_{\theta}(\mu_j) \right] \\
&\stackrel{(d)}{=} - \sum_{k=0}^{K-1} \sum_{\mu_0}^{S \times A} \cdots \sum_{\mu_k}^{S \times A} \sum_{i=0}^k d_0(S_0) \left[\gamma^i \prod_{j=0}^{i-1} \pi_{\theta}(\mu_j) \mathbb{P}(S_{j+1} | \mu_j) \right] \nabla_{\theta} \pi_{\theta}(\mu_i) \left[\prod_{j=i+1}^{k-1} \pi_{\theta}(\mu_j) \mathbb{P}(S_{j+1} | \mu_j) \pi_{\theta}(\mu_k) \gamma^{k-i} R(\mu_k) \right] \\
&\stackrel{(e)}{=} - \sum_{k=0}^{K-1} \sum_{i=0}^k \sum_{S_0}^S d_0(S_0) \sum_{S_i}^S \rho(S_0, S_i, i) \sum_{A_i}^A \nabla_{\theta} \pi_{\theta}(S_i, A_i) \cdot \sum_{\mu_k}^{S \times A} \rho(S_i, S_k, k-i) \cdot \pi_{\theta}(\mu_k) \cdot R(\mu_k) \\
&\stackrel{(f)}{=} - \sum_{S_0}^S d_0(S_0) \sum_S^S \sum_{i=0}^{K-1} \rho(S_0, S, i) \sum_A^A \nabla_{\theta} \pi_{\theta}(S, A) \cdot \left[\sum_{S'}^S \sum_{k=i}^{K-1} \rho(S, S', k-i) \cdot \sum_{A'}^A \pi_{\theta}(S', A') \cdot R(S', A') \right] \\
&\stackrel{(g)}{=} - \sum_{S_0}^S d_0(S_0) \sum_S^S \sum_{i=0}^{\infty} \rho(S_0, S, i) \sum_A^A \nabla_{\theta} \pi_{\theta}(S, A) \cdot Q^{\pi_{\theta}}(S, A) \\
&\stackrel{(h)}{=} - \left[\sum_S^S \sum_{S_0}^S d_0(S_0) \sum_{i=0}^{\infty} \rho(S_0, S, i) \right] \cdot \sum_S^S \frac{\sum_{S_0}^S d_0(S_0) \sum_{i=0}^{\infty} \rho(S_0, S, i)}{\sum_s^S \sum_{S_0}^S d_0(S_0) \sum_{i=0}^{\infty} \rho(S_0, S, i)} \sum_A^A \nabla_{\theta} \pi_{\theta}(S, A) \cdot Q_{\theta}(S, A) \\
&\stackrel{(i)}{\propto} - \sum_S^S d_{\pi_{\theta}}(S) \sum_A^A \nabla_{\theta} \pi_{\theta}(S, A) \cdot Q^{\pi_{\theta}}(S, A) \\
&\stackrel{(j)}{=} - \mathbb{E}_{S \sim d_{\theta}, A \sim \pi_{\theta}(S, \cdot)} [Q^{\pi_{\theta}}(S, A) \nabla_{\theta} \log \pi_{\theta}(S, A)], \tag{26}
\end{aligned}$$

547 where $\rho(S, S', i)$ denotes the probability of state S transfer to S' in i steps.

548 We provide detailed explanations step-by-step:

549 • Equality (a) holds by definition.

550 • In equality (b), using the chain rule, we take derivative of $\nabla_{\theta} [\pi_{\theta}(\mu_0) \cdots \pi_{\theta}(\mu_k)]$ with respect to
551 $\pi_{\theta}(\mu_i)$, $i = 1, \dots, k$.

552 • In equality (c), we plug in L_{μ_0, \dots, μ_k} in (2).

553 • In equality (d), we insert $\mathbb{P}(S_{i+1} | \mu_i) \mathbb{P}(S_{i+1} | \mu_i)$ between $\pi_{\theta}(\mu_i)$ and $\pi_{\theta}(\mu_{i+1})$, $i = 1, \dots, k$.

554 • In equality (e), we split trajectory $\mu_0, \dots, \mu_i, \dots, \mu_k$ into two trajectories μ_0, \dots, μ_i and
555 μ_i, \dots, μ_k . Therefore, we can classify all trajectories μ_0, \dots, μ_k by μ_0, μ_i, μ_k , and i .

556 • In equality (f), we reorganize $\sum_{k=0}^{K-1} \sum_{i=0}^k$ into $\sum_{i=0}^{K-1} \sum_{k=i}^{K-1}$. The former one first traverses the
557 length k of a trajectory, and then traverses the i -th step on it. The latter one first traverses the i -th
558 step of a trajectory, and then traverses the length k of it.

559 • In equality (g), we calculate the limit of (f) when K approaches ∞ .

560 • In equality (h), we normalize $\sum_{S_0}^S d_0(S_0) \sum_{i=0}^{\infty} \rho(S_0, S, i)$ to be a probability distribution.

- 561 • In equality (i), we remove the constant $\sum_S \sum_{S_0}^S d_0(S_0) \sum_{i=0}^{\infty} \rho(S_0, S, i)$ and replace the fraction
 562 with $d_{\pi_\theta}(S)$, the stationary distribution of state S under policy π_θ .
- 563 • In equality (j), we reformulate (i) as expectation.

564

□

565 E.3 PROOF OF THEOREM 4: HAMILTONIAN’S DETERMINISTIC POLICY GRADIENT

566 *Proof.* Let $\eta_\theta(\cdot) : \mathcal{S} \rightarrow \mathcal{A}$ denote a deterministic policy, while we use $\tilde{\pi}_{\theta,\delta}(\mu)$ to represent that a
 567 Gaussian noise (a.k.a, an exploration noise) with standard deviation $\delta > 0$ is added in the exploration
 568 process. In the inference stage, there is no exploration noise, the policy is deterministic, i.e., $\delta = 0$
 569 and $A_k = \eta_\theta(S_k)$.

$$\begin{aligned}
 H'(\theta) &\triangleq -\mathbb{E}_{S_0 \sim d_0, A_0 \sim \tilde{\pi}_{\theta,\delta}} \left[Q^{\tilde{\pi}_{\theta,\delta}}(S_0, A_0) \right] \\
 &= -\mathbb{E}_{S_0, A_k \sim \tilde{\pi}_{\theta,\delta}(S_k, \cdot), S_{k+1} \sim \mathbb{P}(\cdot | S_k, A_k)} \left[\sum_{k=0}^{\infty} \gamma^k \cdot R(S_k, A_k) \right] \\
 &= -\sum_{k=0}^K \mathbb{E}_{S_0, A_k \sim \tilde{\pi}_{\theta,\delta}(S_k, \cdot), S_{k+1} \sim \mathbb{P}(\cdot | S_k, A_k)} \left[\gamma^k \cdot R(S_k, A_k) \right] \\
 &= -\sum_{k=0}^K \sum_{\mu_0}^{\mathcal{S} \times \mathcal{A}} \cdots \sum_{\mu_k}^{\mathcal{S} \times \mathcal{A}} \gamma^k \cdot R(\mu_k) \cdot d_0(S_0) \cdot \tilde{\pi}_{\theta,\delta}(\mu_0) \prod_{i=0}^{k-1} [\mathbb{P}(S_{i+1} | \mu_i) \cdot \tilde{\pi}_{\theta,\delta}(\mu_{i+1})] \quad (27) \\
 &= -\sum_{k=0}^K \sum_{\mu_0}^{\mathcal{S} \times \mathcal{A}} \cdots \sum_{\mu_k}^{\mathcal{S} \times \mathcal{A}} \left[\gamma^k \cdot R(\mu_k) \cdot d_0(S_0) \cdot \prod_{i=0}^{k-1} \mathbb{P}(S_{i+1} | \mu_i) \right] \cdot \tilde{\pi}_{\theta,\delta}(\mu_0) \cdots \tilde{\pi}_{\theta,\delta}(\mu_k) \\
 &= -\sum_{k=0}^K \sum_{\mu_0}^{\mathcal{S} \times \mathcal{A}} \cdots \sum_{\mu_k}^{\mathcal{S} \times \mathcal{A}} L_{\mu_0, \dots, \mu_k} \cdot \tilde{\pi}_{\theta,\delta}(\mu_0) \cdots \tilde{\pi}_{\theta,\delta}(\mu_k),
 \end{aligned}$$

570 where $K \rightarrow \infty$, and

$$L_{\mu_0, \dots, \mu_k} = \gamma^k \cdot R(\mu_k) \cdot d_0(S_0) \cdot \prod_{i=0}^{k-1} \mathbb{P}(S_{i+1} | \mu_i). \quad (28)$$

571

□

572 E.4 PROOF OF COROLLARY 2

Proof.

$$\begin{aligned}
\nabla_{\theta} H'(\pi_{\theta}) &= - \sum_{k=0}^K \sum_{\mu_0}^{S \times A} \cdots \sum_{\mu_k}^{S \times A} (L_{\mu_0, \dots, \mu_k} \cdot \nabla_{\theta} [\tilde{\pi}_{\theta}(\mu_0) \cdots \tilde{\pi}_{\theta}(\mu_k)] + \nabla_{\theta} L_{\mu_0, \dots, \mu_k} \cdot \tilde{\pi}_{\theta}(\mu_0) \cdots \tilde{\pi}_{\theta}(\mu_k)) \\
&= - \sum_{k=0}^K \sum_{\mu_0}^{S \times A} \cdots \sum_{\mu_k}^{S \times A} [\tilde{\pi}_{\theta}(\mu_0) \cdots \tilde{\pi}_{\theta}(\mu_k)] \cdot \nabla_{\theta} L_{\mu_0, \dots, \mu_k} \\
&= - \sum_{k=0}^K \sum_{\mu_0}^{S \times A} \cdots \sum_{\mu_k}^{S \times A} \nabla_{\theta} \left[\gamma^k \cdot R(\mu_k) \cdot d_0(S_0) \cdot \prod_{i=0}^{k-1} \mathbb{P}(S_{i+1} | \mu_i) \right] \\
&= - \sum_{k=0}^K \sum_{\mu_0}^{S \times A} \cdots \sum_{\mu_k}^{S \times A} \nabla_A \left[\gamma^k \cdot R(\mu_k) \cdot d_0(S_0) \cdot \prod_{i=0}^{k-1} \mathbb{P}(S_{i+1} | \mu_i) \right] \nabla_{\theta} \eta_{\theta}(S) \\
&= - \sum_{S_0}^{\mathcal{S}} d_0(S_0) \nabla_A \mathbb{E}_{S_{t+1} \sim \mathbb{P}(\cdot | S_t, A_t)} \left[\sum_{t=0}^{\infty} \gamma^k R(S_t, A_t) \right] \cdot \nabla_{\theta} \eta_{\theta}(S) \\
&= - \sum_{S_0}^{\mathcal{S}} d_0(S_0) \nabla_A Q(S_0, A_0) \cdot \nabla_{\theta} \eta_{\theta}(S) \\
&= - \mathbb{E}_{S_0 \sim d_0(\cdot)} [\nabla_A Q(S_0, A_0) \cdot \nabla_{\theta} \eta_{\theta}(S)]
\end{aligned} \tag{29}$$

573 where $\mu_k = (S_k, A_k)$, $S_0 \sim d_0(\cdot)$, $A_k \sim \pi_{\theta}(S_k, \cdot)$, $S_{k+1} \sim \mathbb{P}(\cdot | S_k, A_k)$, for $k = 0 \cdots K$. \square

574 E.5 PROOF OF COROLLARY 3

575 *Proof.* In Corollary 2 and Corollary 1, we have

$$\begin{aligned}
\nabla_{\theta} H'(\theta) &= -\nabla_{\theta} J'(\theta), \\
\nabla_{\theta} H(\theta) &= -\nabla_{\theta} J(\theta),
\end{aligned} \tag{30}$$

576 when $K \rightarrow \infty$.

577 [43] proved that

$$\nabla_{\theta} J'(\theta) = \lim_{\delta \rightarrow 0} \nabla_{\theta} J(\theta), \tag{31}$$

578 where δ is the standard deviation of the Gaussian noise of stochastic policy π_{θ} .

579 Therefore,

$$\nabla_{\theta} H'(\theta) = \lim_{\delta \rightarrow 0} \nabla_{\theta} H(\theta) \tag{32}$$

580 \square

581 F VARIANCE REDUCTION WITH QUANTUM K-SPIN HAMILTONIAN EQUATION

582 F.1 MONTE CARLO ESTIMATOR OF QUANTUM K-SPIN HAMILTONIAN EQUATION

583 **Monte Carlo Estimator** [38]: Consider a general probabilistic objective \mathcal{F} of the form:

$$\mathcal{F} \triangleq \mathbb{E}_{p(\mathbf{x};\theta)}[f(\mathbf{x};\phi)], \quad (33)$$

584 in which a function f of an input variable \mathbf{x} with *structural parameters* ϕ is evaluated on average
585 with respect to an input distribution $p(\mathbf{x};\theta)$ with *distributional parameters* θ .586 A Monte Carlo method evaluates the function by first drawing independent samples $\hat{\mathbf{x}}^{(1)}, \dots, \hat{\mathbf{x}}^{(N)}$
587 from the distribution $p(\mathbf{x};\theta)$, and then computing the average:

$$\hat{\mathcal{F}}_N = \frac{1}{N} \sum_{i=1}^N f(\hat{\mathbf{x}}^{(i)}), \quad \text{where } \hat{\mathbf{x}}^{(i)} \sim p(\mathbf{x};\theta) \text{ for } i = 1, \dots, N. \quad (34)$$

588 The Monte Carlo estimator for (15) is

$$\hat{\mathcal{J}}(\theta) = \frac{1}{N} \sum_{i=1}^N R(\tau^{(i)}), \quad \text{where } \tau^{(i)} \sim P(\tau^{(i)}|\pi_\theta) \text{ for } i = 1, \dots, N, \quad (35)$$

589 and

590

$$P(\tau^{(i)}|\pi_\theta) = d_0(s_0^{(i)}) \cdot \prod_{k=0}^T \mathbb{P}(s_{k+1}^{(i)}|s_k^{(i)}, a_k^{(i)}) \pi_\theta(a_k^{(i)}|s_k^{(i)}). \quad (36)$$

591 The Monte Carlo estimator for (3) is

$$\hat{H}(\theta) = \frac{1}{N'} \sum_{i=1}^{N'} \sum_{k=0}^{K-1} L_{\mu_0^{(i)}, \dots, \mu_k^{(i)}}, \quad \text{for } i = 1, \dots, N', \quad (37)$$

592 and

$$L_{\mu_0^{(i)}, \dots, \mu_k^{(i)}} = \gamma^k \cdot R(\mu_k^{(i)}) \cdot d_0(s_0^{(i)}) \cdot \prod_{\ell=0}^{k-1} \mathbb{P}(s_{\ell+1}^{(i)}|\mu_\ell^{(i)}). \quad (38)$$

593 **Remark:** The above two Monte Carlo estimators are quite different in the simulation process. (36)
594 samples a random trajectory by following an environment’s stochastic transition and a policy. In
595 contrast, (38) measures a random path’s discounted reward (the “energy”) without following any
596 policy, and the Hamiltonian equation (3) combinatorially enumerates all possible paths of length K
597 over the state-action space. In other words, the simulation process of the Hamiltonian term does not
598 rely on any policy. Therefore, the Hamiltonian term is a suitable regularizer for both on-policy and
599 off-policy algorithms.

600 This fundamental difference is due to the Ising model in (1), which combinatorially enumerates all
601 paths and separates the environment and the policy.

602 F.2 VARIANCE REDUCTION

603 For the general function in (33), one simple but effective variance reduction technique is to subtract a
604 baseline term as follows:

$$\mathbb{E}_{p(\mathbf{x};\theta)} [(f(\mathbf{x}) - \beta) \nabla_\theta \log p(\mathbf{x};\theta)], \quad (39)$$

605 where β is the baseline term.606 **Our reasoning logic:**

607 1). We first briefly describe a high-level idea [22] that adding a baseline term, like the proposed
608 H-term, will help reduce the gradient variance.

609 **2). We sketch the steps to show how the proposed H-term will mathematically reduce the gradient**
 610 **variance, following the framework in Section 5.2 of [23].**

611 **High-level IDEA.** One generic approach to reduce the variance of Monte Carlo estimates is to use an
 612 additive control variate. Suppose we wish to estimate the integral of the function $f : \mathcal{X} \rightarrow \mathbb{R}$, and we
 613 know the value of the integral of another function on the same space $\phi : \mathcal{X} \rightarrow \mathbb{R}$.

614 We have

$$\int_{\mathcal{X}} f(x) = \int_{\mathcal{X}} (f(x) - \phi(x)) + \int_{\mathcal{X}} \phi(x), \quad (40)$$

615 and the integral of $f(x) - \phi(x)$ can be estimated. If $\phi(x) = f(x)$, meaning that , then we have
 616 managed to reduce our variance to zero [22]. More generally,

$$\text{Var}(f - \phi) = \text{Var}(f) - 2\text{Cov}(f, \phi) + \text{Var}(\phi). \quad (41)$$

617 If ϕ and f are strongly correlated, so that the covariance term on the right hand side is greater than
 618 the variance of ϕ , i.e., $-2\text{Cov}(f, \phi) + \text{Var}(\phi) \leq 0$. then a variance reduction has been made over the
 619 original estimation problem [22], i.e., $\text{Var}(f - \phi) \leq \text{Var}(f)$.

620 **Our reasoning.** Then, we present our reasoning.

621 Note that the gradient of the new objective function of the actor network in (8) consists of two
 622 components, namely $\nabla_{\theta} J(\theta)$ and $\nabla_{\theta} H(\theta)$. Here, we consider

$$\nabla_{\theta} J(\theta) - \lambda \nabla_{\theta} H(\theta), \quad \text{where } \lambda > 0 \text{ is a temperature parameter,} \quad (42)$$

623 where $\nabla_{\theta} J(\theta)$ in (21) is the above function $f(\cdot)$ and $\lambda \nabla_{\theta} H(\theta)$ in (4) is the above function $\phi(\cdot)$.

624 The Hamiltonian stochastic gradient in (4) has the optimal value

$$\nabla_{\theta} H^*(\theta) = - \lim_{K \rightarrow \infty} \mathbb{E}_{\mu_0, \dots, \mu_{K-1}} \left[\sum_{k=0}^{K-1} \gamma^k \cdot R(\mu_k) \cdot \nabla_{\theta} \log(\pi_{\theta}(\mu_0) \cdot \pi_{\theta}(\mu_1) \cdots \pi_{\theta}(\mu_k)) \right]. \quad (43)$$

625 According to Theorem 8 of [22] that is proved via (41), we have

$$\begin{aligned} \text{Var}[\nabla_{\theta} J(\theta) - \lambda \nabla_{\theta} H^*(\theta)] &= \text{Var}[\nabla_{\theta} J(\theta)] - \frac{1}{\lambda} \mathbb{E}_{s \sim d_{\theta}, a \sim \pi_{\theta}} \left[\frac{(\mathbb{E}_{s \sim d_{\theta}, a \sim \pi_{\theta}} [(\nabla_{\theta} \log \pi_{\theta}(s, a))^2 \nabla_{\theta} J(\theta)])^2}{\mathbb{E}_{s \sim d_{\theta}, a \sim \pi_{\theta}} [(\nabla_{\theta} \log \pi_{\theta}(s, a))^2]} \right] \\ &\leq \text{Var}[\nabla_{\theta} J(\theta)], \end{aligned} \quad (44)$$

626 where the second term is positive and

$$\nabla_{\theta} H^*(\theta) = \frac{\mathbb{E}_{s \sim d_{\theta}, a \sim \pi_{\theta}} [(\nabla_{\theta} \log \pi_{\theta}(s, a))^2 \nabla_{\theta} J(\theta)]}{\mathbb{E}_{s \sim d_{\theta}, a \sim \pi_{\theta}} [(\nabla_{\theta} \log \pi_{\theta}(s, a))^2]}. \quad (45)$$

627 In Alg. 1 and Alg. 2, we used a general H-term $\nabla_{\theta} H(\theta)$, not the optimal one in (43). Next, we
 628 provide a general characterization for this case.

629 According to Theorem 10 of [22], we have

$$\begin{aligned} &\text{Var}[\nabla_{\theta} J(\theta) - \lambda \nabla_{\theta} H(\theta)] - \text{Var}[\nabla_{\theta} J(\theta) - \lambda \nabla_{\theta} H^*(\theta)] \\ &= \lambda^2 \mathbb{E}_{s \sim d_{\theta}, a \sim \pi_{\theta}} [(\nabla_{\theta} \log \pi_{\theta}(s, a))^2 (\nabla_{\theta} H(\theta) - \nabla_{\theta} H^*(\theta))^2] \end{aligned} \quad (46)$$

630 Assume Lipschitz continuity of the gradient $\nabla_{\theta} H(\theta)$ such that

$$\|\nabla_{\theta} H(\theta) - \nabla_{\theta} H^*(\theta)\|_2 \leq 2L(H(\theta) - H^*(\theta)) \leq 2L\epsilon, \quad (47)$$

631 given $K \geq \log_{\gamma} \epsilon$ with $L > 0, \epsilon > 0$, as pointed out in the end of Section 3.2.

632 Therefore, combining (46), (48) with (48), we obtain that

$$\begin{aligned}
\text{Var}[\nabla_{\theta}J(\theta) - \lambda \nabla_{\theta}H(\theta)] &= \text{Var}[\nabla_{\theta}J(\theta)] - \frac{1}{\lambda} \mathbb{E}_{s \sim d_{\theta}, a \sim \pi_{\theta}} \left[\frac{(\mathbb{E}_{s \sim d_{\theta}, a \sim \pi_{\theta}} [(\nabla_{\theta} \log \pi_{\theta}(s, a))^2 \nabla_{\theta}J(\theta)])^2}{\mathbb{E}_{s \sim d_{\theta}, a \sim \pi_{\theta}} [(\nabla_{\theta} \log \pi_{\theta}(s, a))^2]} \right] \\
&\quad + \lambda^2 (2L\epsilon)^2 \mathbb{E}_{s \sim d_{\theta}, a \sim \pi_{\theta}} [(\nabla_{\theta} \log \pi_{\theta}(s, a))^2] \\
&\leq \text{Var}[\nabla_{\theta}J(\theta)],
\end{aligned} \tag{48}$$

633 when both $|\nabla_{\theta} \log \pi_{\theta}(s, a)|$ and $|\nabla_{\theta}J(\theta)|$ are upper bounded, e.g., $|\nabla_{\theta} \log \pi_{\theta}(s, a)| < C_1$ and
634 $|\nabla_{\theta}J(\theta)| < C_2$; and we set ϵ, λ properly such that

$$\begin{aligned}
-\frac{1}{\lambda} C_2^2 + 4\lambda^2 L^2 \epsilon^2 C_1^2 &< 0 \\
\lambda^3 \epsilon^2 &< \frac{C_2^2}{4L^2 C_1^2},
\end{aligned} \tag{49}$$

635 which can be easily satisfied by properly selecting λ and $K \geq \log_{\gamma} \epsilon$.

636 **Conclusion:**

637 To sum up, we show that it is easy to achieve $\text{Var}[\nabla_{\theta}J(\theta) - \lambda \nabla_{\theta}H(\theta)] \leq \text{Var}[\nabla_{\theta}J(\theta)]$, which
638 means adding the H-term can lead to smaller variance than that of the conventional gradient.

639 G ACTOR-CRITIC ALGORITHMS FOR DEEP REINFORCEMENT LEARNING

640 The gradient of (15) is [44]

$$\nabla_{\theta} J(\theta) \triangleq \sum_S d_{S,\theta}(S) \sum_A Q_{\theta}(S, A) \nabla_{\theta} \pi_{\theta}(S, A). \quad (50)$$

641 Since Q_{θ} in (50) is unknown [49] (the stationary distribution d_{θ} is unknown), one can plug in a critic
642 network with parameter ϕ as an estimator of Q_{θ} and obtain

$$\nabla_{\theta}^{\phi} J(\theta, \phi) = \sum_S d_{S,\theta}(S) \sum_A Q_{\phi}(S, A) \nabla_{\theta} \pi_{\theta}(S, A), \quad (51)$$

643 where $d_{S,\theta} \in \mathbb{R}_+^{|S||A| \times 1}$ denotes the stationary distribution over the states instead of state-action
644 pairs.

645 (51) is a bi-level optimization problem [9], and a natural solution is an iterative algorithm that
646 alternates between estimating Q_{ϕ} with parameter ϕ and improving policy π_{θ} with parameter θ .
647 Therefore, a family of actor-critic algorithms are proposed with following objective functions:

$$\begin{cases} \text{Actor : } \max_{\theta} J_{\pi}(\theta, \phi) = (1 - \gamma) \mathbb{E}_{S_0 \sim d_0, A_0 \sim \pi_{\theta}(S_0, \cdot)} [Q_{\phi}(S_0, A_0)] \\ \text{Critic : } \max_{\phi} J_Q(\theta, \phi) = \frac{1}{2} \mathbb{E}_{S \sim d_{\theta}(\cdot), A \sim \pi_{\theta}(S, \cdot)} [(Q_{\phi}(S, A) - y(S, A))^2]. \end{cases} \quad (52)$$

648 The gradient of (52) can be estimated as follows

$$\begin{aligned} \nabla_{\theta} \hat{J}_{\pi}(\theta, \phi) &= \frac{1}{N} \sum_{i=1}^N Q_{\phi}(\mu) \cdot \nabla_{\theta} \log \pi_{\theta}(\mu) \\ \nabla_{\phi} \hat{J}_Q(\theta, \phi) &= \frac{1}{N} \sum_{i=1}^N [Q_{\phi}(S, A) - y(S, A)] \cdot \nabla_{\phi} Q_{\phi}(S, A) \end{aligned} \quad (53)$$

649 The parameters ϕ and θ are updated as follows:

$$\begin{cases} \text{Actor : } \theta \leftarrow \theta + \alpha \nabla_{\theta}^{\phi} \hat{J}_{\pi}, \\ \text{Critic : } \phi \leftarrow \phi - \alpha \nabla_{\phi} \hat{J}_Q. \end{cases} \quad (54)$$

Algorithm 2 Stationary Actor-Critic Algorithm with H-term (Deterministic Version)

```

1: Input: learning rate  $\alpha$ , temperature  $\lambda$ , look-ahead step  $K$ , and parameters  $\delta, M, T, G, B, B'$ 
2: Initialize actor network  $\eta$  and critic network  $Q$  with parameters  $\theta, \phi$ , and replay buffers  $\mathcal{D}_1, \mathcal{D}_2$ 
3: for episode = 1,  $\dots$ ,  $M$  do
4:   Initialize state  $s_0$ 
5:   for  $t = 0, \dots, T - 1$  do
6:     Take action  $a_t = \eta_\theta(s_t) + \epsilon$ , where  $\epsilon \sim \mathcal{N}(0, \delta^2)$ 
7:     Execute action  $a_t$ , receive reward  $r_t$ , and observe new state  $s_{t+1}$ 
8:     Store a transition  $(s_t, a_t, r_t, s_{t+1})$  in  $\mathcal{D}_1$ 
9:   end
10:  Store a trajectory  $\tau$  of length  $T$  in  $\mathcal{D}_2$ 
11:  for  $g = 1, \dots, G$  do
12:    Randomly sample a mini-batch of  $B$  transitions  $\{(s_i, a_i, r_i, s_{i+1})\}_{i=1}^B$  from  $\mathcal{D}_1$ 
13:    Randomly sample a mini-batch of  $B'$  trajectories (of length  $K$ )  $\{\tau_j\}_{j=1}^{B'}$  from  $\mathcal{D}_2$ 
14:    Update critic network using a conventional method
15:    Update actor network as  $\theta \leftarrow \theta + \alpha \left( \nabla_\theta \hat{J}'(\theta) - \lambda \nabla_\theta \hat{H}'(\theta) \right)$ .
16:  end
17: end

```

650 H DETERMINISTIC POLICY GRADIENT ALGORITHM WITH H-TERM

651 For completeness, we present the details of the deterministic actor-critic algorithm with H-term.

652 We apply the proposed Hamiltonian equation (3) to regularize the actor network. Specifically, $H'(\theta)$
653 in (3) is added to the actor's objective with weight $\lambda > 0$. The objective functions of actor and critic
654 networks become:

$$\begin{cases} \text{Actor : } \max_{\theta} J'_\pi(\theta, \phi) = (1 - \gamma) \mathbb{E}_{S_0 \sim d_0, A_0 = \eta_\theta(S_0)} [Q_\phi(S_0, A_0)] - \lambda H'(\theta), \\ \text{Critic : } \min_{\phi} J_Q(\theta, \phi) = \frac{1}{2} \mathbb{E}_{S \sim d_\theta(\cdot), A = \eta_\theta(S)} [(Q_\phi(S, A) - y(S, A))^2]. \end{cases} \quad (55)$$

655 The gradient of (55) is

$$\nabla_\theta J'_\pi(\theta, \phi) = (1 - \gamma) \sum_S d_{S, \theta}(S) \nabla_A Q_\phi(S, A) \cdot \nabla_\theta \eta_\theta(S) - \lambda \nabla_\theta H'(\theta), \quad (56)$$

656

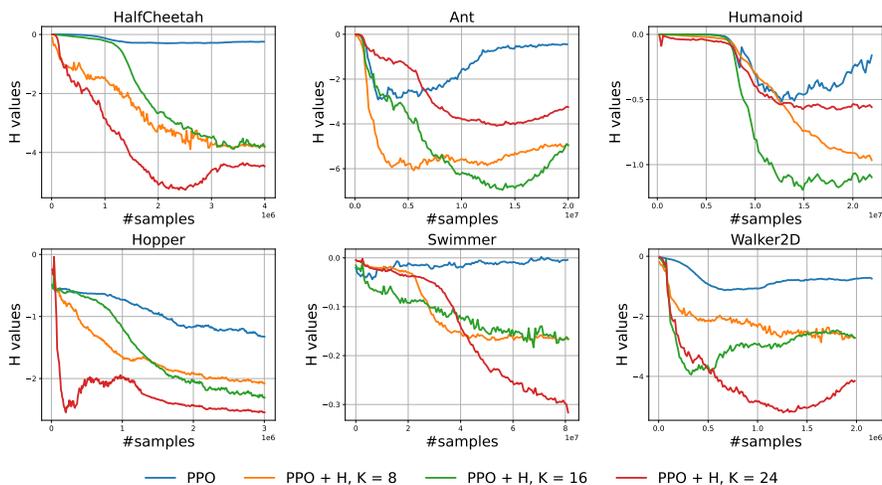
$$\nabla_\phi J_Q(\theta, \phi) = \sum_S d_{S, \theta}(S) \cdot [Q_\phi(S, A) - y(S, A)] \cdot \nabla_\phi Q_\phi(S, A)|_{A = \eta_\theta(S)}. \quad (57)$$

657 To estimate $\nabla_\theta H'(\theta)$, the Monte Carlo gradient estimator in (6) is used. Therefore, (56) and (57) can
658 be estimated as follows:

$$\nabla_\theta \hat{J}'_\pi(\theta, \phi) = \frac{1}{N} \sum_{i=1}^N [\nabla_A Q_\phi(S, A)|_{A = \eta_\theta(S)} \nabla_\theta \eta_\theta(S)] - \frac{1}{N'} \sum_{i=1}^{N'} \left[\lambda \sum_{k=0}^K \gamma^k R(\mu_k) \nabla_\theta \log [\tilde{\pi}_\theta(\mu_0) \cdots \tilde{\pi}_\theta(\mu_k)] \right], \quad (58)$$

659

$$\nabla_\phi \hat{J}_Q(\theta, \phi) = \frac{1}{N} \sum_{i=1}^N [Q_\phi(S, A) - y(S, A)] \cdot \nabla_\phi Q_\phi(S, A)|_{A = \eta_\theta(S)}. \quad (59)$$

Figure 7: H values during the training process.

660 I EXPERIMENTS: HYPERPARAMETERS AND MORE RESULTS

661 I.1 HYPERPARAMETERS IN EXPERIMENTS

Table 5: Hyperparameters used for the PPO and PPO + H in MuJoCo tasks

Parameters	Values
Optimizer	Adam
Learning rate	$3 \cdot 10^{-4}$
Discount (γ)	0.99
GAE parameter	0.95
Number of hidden layers for all networks	3
Number of hidden units per layer	256
Mini-batch size	32
Importance rate of H-term (λ)	2^{-3}
Truncation step of H-term (K)	16

Table 6: Hyperparameters used for the DDPG and DDPG + H in MuJoCo tasks

Parameters	Values
Optimizer	Adam
Learning rate	$5 \cdot 10^{-4}$
Target Update Rate (τ)	10^{-3}
Discount (γ)	0.995
Replay buffer size	10^6
Number of hidden layers for all networks	3
Number of hidden units per layer	256
Batch size	64
Importance rate of H-term (λ)	2^{-3}
Truncation step of H-term (K)	16

662 I.2 MORE RESULTS

663 Fig. 7 shows the H-value (average over 20 runs) during the training process, which verified that the
 664 trained agents have converged to policies with small H-values.

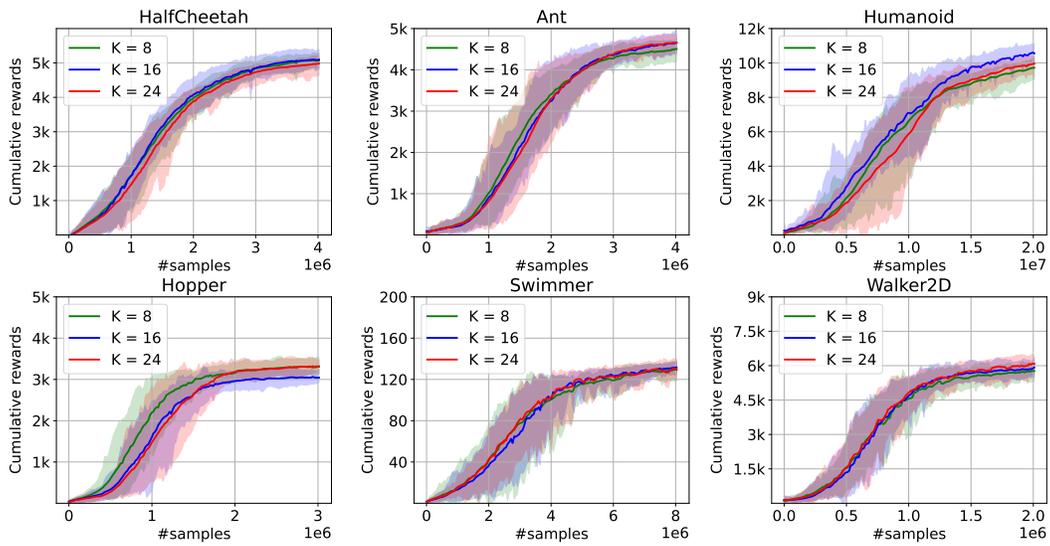


Figure 8: For the proposed PPO+H algorithm, the performance with different K values.

665 Fig. 8 shows more performance of the PPO+H algorithm, for $K = 8, 16, 24$. We run each experiment
 666 with 20 random seeds and each run we test 100 episodes.

667 To verify the hypothesis that smaller replay buffer hurts the performance, we rerun the trials of
 668 $K = 8, 16$ with a replay buffer size 800.

Algorithm 3 Hamiltonian Policy Network

```

1: Input: learning rate  $\alpha$ , look-ahead step  $K$ , and parameters  $M, T, G, B$ 
2: Initialize policy network with parameters  $\theta$ , and replay buffer  $\mathcal{D}$ 
3: for episode = 1,  $\dots$ ,  $M$  do
4:   Initialize state  $s_0$ 
5:   for  $t = 0, \dots, T - 1$  do
6:     Select action  $a_t \sim \pi_\theta(\cdot|s_t)$ 
7:     Execute action  $a_t$ , receive reward  $r_t$ , and observe new state  $s_{t+1}$ 
8:   end
9:   Store a trajectory  $\tau$  of length  $T$  in  $\mathcal{D}$ 
10:  for  $g = 1, \dots, G$  do
11:    Randomly sample a mini-batch of  $B$  trajectories (of length  $K$ )  $\{\tau_j\}_{j=1}^B$  from  $\mathcal{D}$ 
12:    Update policy network as  $\theta \leftarrow \theta - \alpha \nabla_\theta \hat{H}(\theta)$ .
13:  end
14: end

```

669 **J HAMILTONIAN POLICY NETWORK**670 **J.1 HAMILTONIAN POLICY NETWORK**

671 Since Hamiltonian equation in (3) is a functional of policy π_θ , a natural question would be: can
672 we use the Hamiltonian equation replace existing Bellman’s equation (16) or the policy gradient’s
673 objective function (15)?

674 As a verification, we test the capability of Hamiltonian equation in (3) as a loss function to train a
675 policy network. The algorithm is first given as follows.

676 In Alg. 3, an agent interacts with an environment and updates its policy network. The algorithm has
677 M episodes and each episode consists of a (Monte Carlo) simulation process and a learning process
678 (gradient estimation) as follows :

- 679 • During the (Monte Carlo) simulation process (lines 5-9 of Alg. 3), an agent takes action a_t
680 according to a policy $\pi_\theta(\cdot|s_t)$, $t = 0, \dots, T - 1$, generating a trajectory of T steps/transitions.
681 Then, the full trajectory $\tau = (s_0, a_0, r_0, s_1, \dots, s_{T-1}, a_{T-1}, r_{T-1}, s_T)$ is stored in replay buffer
682 \mathcal{D} .
- 683 • During the learning process ($G \geq 1$ updates in one episode) (lines 10-12 of Alg. 1), a mini-batch of
684 B trajectories (of length K) $\{\tau_j = (s_0^j, a_0^j, r_0^j, s_1^j, \dots, s_{K-1}^j, a_{K-1}^j, r_{K-1}^j, s_K^j)\}_{j=1}^B$ are sampled
685 from \mathcal{D} , respectively. The policy network is updated by a Monte Carlo gradient estimator over B
686 trajectories.

687 **Implementation of replay buffer \mathcal{D} .** After a full trajectory τ of length T is generated, it is partitioned
688 into $T - K + 1$ trajectories of length K . We rank them according to the cumulative reward and
689 store the top portion, say 80%, into a new replay buffer \mathcal{D} (line 9 of Alg. 3). We randomly sample a
690 mini-batch of B trajectories from \mathcal{D} (line 11 of Alg. 3) to compute the H-term.

691 **J.2 FROZENLAKE TASK**

692 **Environment:** Frozenlake 8×8 , a game in OpenAI Gym.

693 **Rules:** As shown in Fig. 9 (left), the Frozenlake task has 8×8 states with 4 optional actions to move
694 around. The agent needs to go from the start point and find the way to the destination in limited steps.
695 There are 8 holes which can cause the agent to fail the game.

696 **Experiment settings:** We take Deep Q-learning (DQN) [37] as our baseline and use the implementa-
697 tion from the ElegantRL library. We use a 4-layer fully connected neural network as the deep policy
698 network both in DQN and DHN. We use the Adam optimizer with a learning rate 1×10^{-3} and a
699 batch size 100.

700 **Evaluation:** We evaluate the performance of policy by computing the success rate, in which we use
701 50 agents to walk 100 steps and compute the rates of agents who successfully arrive the destination.

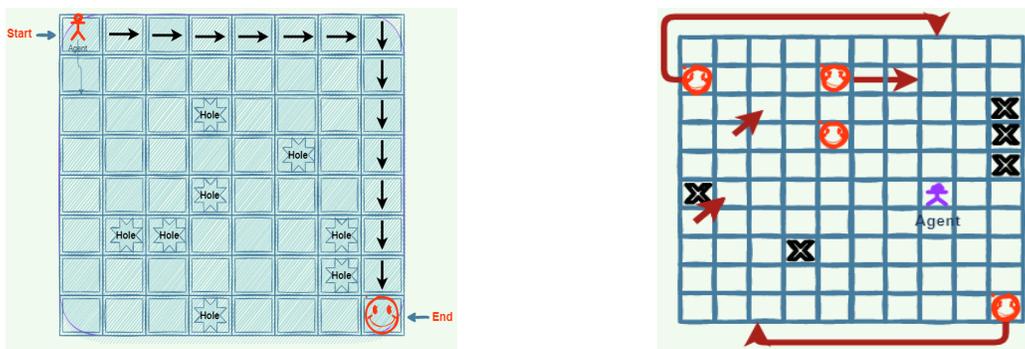


Figure 9: The Frozenlake task (left) and Gridworld task (right).

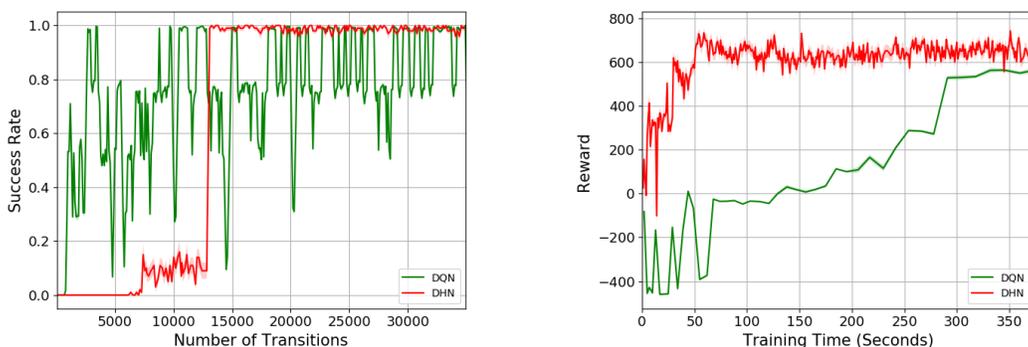


Figure 10: Comparison between the DQN and DHN algorithms. The Frozenlake task (left) and Gridworld task (right).

702 **Results for the Frozenlake task:** Fig. 10 (left) shows the success rate of agents with increasing the
 703 number of transitions learned by the network. compared with DQN, DHN has a more stable training
 704 process. It is easy for DQN to quickly obtain a good policy to win the game. But with increasing the
 705 number of transitions fed to the network, the performance of DQN shows a large and frequent shock
 706 while the performance of DHN shows the strong stability.

707 J.3 GRIDWORLD TASK

708 **Environment:** a Gridworld of size 10×10 , a game available in our code.

709 **Rules:** As shown in the Fig. 9 (right), the Gridworld has 10×10 states with 4 optional actions to
 710 move around. The agent will initialize at a random locations and it needs to find the smiley as many
 711 as possible which has 10 reward in turn. It should be noted that there are some endpoints which may
 712 cause the agent game over and some transfer-points which transfer the agent to certain location.

713 **Experiment settings and evaluation:** Both the experiment settings and evaluation method are the
 714 same with that on Frozenlake 8×8 game.

715 **Results for the Gridworld task:** Fig. 10 (right) shows the mean reward obtained by the agents with
 716 increasing the training time. Compared with DQN, DHN has a faster training process. It only needs
 717 massive random parallel samples of trajectories and do not need any policy for guided sampling while
 718 DQN needs guided exploration in the training process which costs a large time consumption.

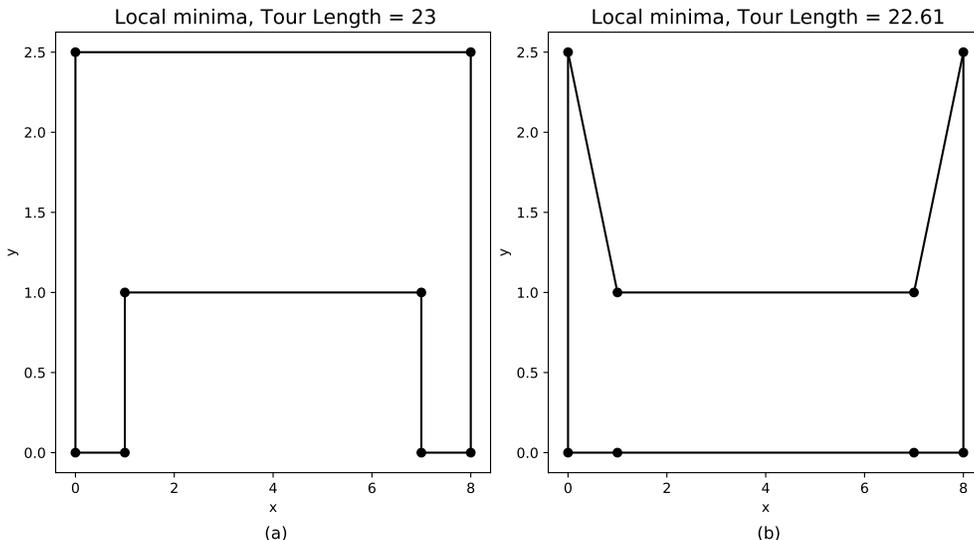


Figure 11: TSP problem has two local minimas in this case.

719 **K COMBINATORIAL OPTIMIZATION PROBLEMS**

720 **K.1 PROBLEM FORMULATION AND MDP FORMULATION**

- 721 • **Graph maxcut:** Given a graph $G = (V, E)$, where V is the set of nodes and E the set of edges,
722 find a subset $S \subseteq V$ that maximizes the weight of the cut-set $\sum_{u \in S, v \in V \setminus S, (u,v) \in E} w(u, v)$.
- 723 • **Traveling salesman problems (TSP):** Given a fully connected graph $G = (V, E)$, find a tour J
724 that minimizes the edge weights $\sum_{e \in J} w(e)$. A tour starts and ends at a specific node after having
725 visited each node exactly once.
- 726 • **MDP formulation:** The MDP formulations of graph maxcut and TSP are given in Table 7. Note,
727 J is a partially ordered set, (J, v) means add node v to the end of J .

728 **K.2 EXISTENCE OF MULTIPLE FIXED POINTS**

- 729 • Travelling salesman problem (TSP): the case of 8 cities (Fig. 11) has 2 local minimas.
- 730 • Graph maxcut: a fully connected graph of 20 nodes (Appx. K) has 390 local minimas.

731 **K.3 EXPERIMENTAL EVALUATION**

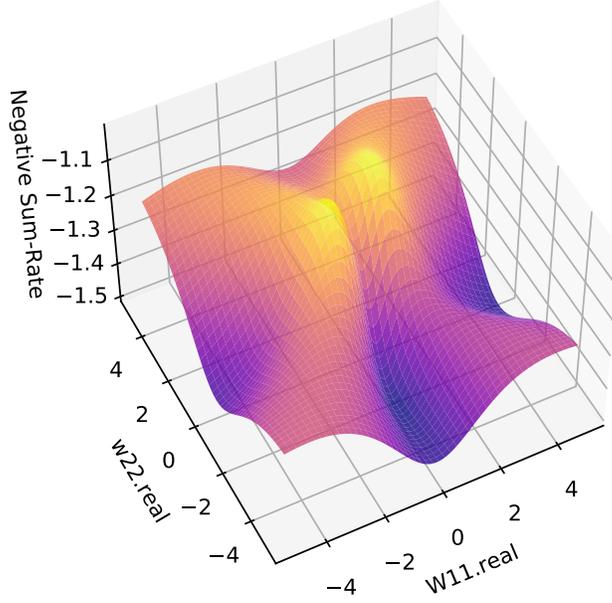
- 732 • **Environments (tasks).** We evaluate the proposed algorithms on graph maxcut and TSP tasks with
733 number of nodes $N = 100$. We generate instances for graph maxcut and TSP seperately. For the
734 graph maxcut problem, we generate 1000 Erdős-Renyi (ER) graphs as test instances. For the TSP
735 problem, we use the same test instances in [31].
- 736 • **Compared algorithms.** To evaluate policy gradient algorithm, we choose REINFORCE [45].
737 Since the H-term is compatible with existing variance reduction techniques, we implement the RE-
738 INFORCE algorithm with baseline reduction. For a fair comparison, we keep the hyperparameters
739 (listed in Appx. I) the same and make sure that the obtained results reproduce existing benchmark
740 tests.
- 741 • **Performance metrics.** We employ the metric: Frequency vs. Approximation Ratio. We run each
742 experiment with 20 random seeds and in each run we test 100 episodes. The approximation ratio ϵ
743 is defined as

$$\epsilon = \max\left(\frac{Optimal}{Obj}, \frac{Obj}{Optimal}\right), \tag{60}$$

744 where *Optimal* is the optimal objective value by Branch and Bound [32], *Obj* is the objective
745 value of the best solution obtained from RL.

Table 7: MDP formulation of combinatorial optimization: Graph maxcut and TSP

Problem	State	Action	Transition	Reward	Termination
Graph maxcut	S	select $v \in V \setminus S$	$S' \leftarrow S \cup \{v\}$	$\text{cut}(S') - \text{cut}(S)$	$\forall S', \text{cut}(S') < \text{cut}(S)$
TSP	J	select $v \in V \setminus J$	$J' \leftarrow (J, v)$	$\text{cost}(J') - \text{cost}(J)$	$\forall v \in V, v \in J$

Figure 12: Geometry of sum-rate ($K = 2, N = 2, P = 10$) with respect to the real part of W_{11} and W_{22} . It has three local minimas.

746 L NON-CONVEX OPTIMIZATION PROBLEMS

- 747 • MIMO beamforming [7]: Given a channel $H = [\mathbf{h}_1, \dots, \mathbf{h}_M] \in \mathbb{C}^{N \times M}$, and a power constraint
 748 $P \in \mathbb{R}$, find a beamformer $W = [\mathbf{w}_1, \dots, \mathbf{w}_M] \in \mathbb{C}^{N \times M}$ that maximizes the summation of rates:
 749

$$\begin{aligned} & \max_{\mathbf{W}} \sum_{m=1}^M \log_2(1 + \text{SINR}_m) \\ & \text{s.t. } \text{Tr}(\mathbf{W}\mathbf{W}^H) \leq P, \end{aligned} \quad (61)$$

where $\text{SINR}_m = \frac{|\mathbf{h}_m^H \mathbf{w}_m|^2}{\sum_{j=1, j \neq m}^M |\mathbf{h}_m^H \mathbf{w}_j|^2 + \sigma^2}$.

- 750 • Non-convex deep learning classifier: Consider a 2-layer classifier, using ReLU as activation layer
 751 and softmax as output layer. The task is to find weights W_1, W_2 and biases b_1, b_2 that maximize the
 752 cross entropy loss:

$$\min_{W_1, W_2, b_1, b_2} -\frac{1}{n} \sum_{i=1}^n \log \left(\frac{\exp \left((W_2 \max(W_1 \mathbf{x}_i + \mathbf{b}_1, 0) + \mathbf{b}_2)_{y_i} \right)}{\sum_j \exp \left((W_2 \max(W_1 \mathbf{x}_i + \mathbf{b}_1, 0) + \mathbf{b}_2)_j \right)} \right), \quad (62)$$

- 753 where W_i and b_i are the weight and bias of i -th layer, n is the batch size, x_i is the input, and y_i
 754 is the target class.