

Knowledge-Assistant Deep Reinforcement Learning for Multi-Agent Region Protection

Siqing Sun and Tianbo Li

Since the rewards are newly designed, the convergence of the critical neural network needs to be proven. Once the error of the critical neural network converges, the Bellman Optimality Equation can be applied, ensuring the feasibility and convergence of the solution to Problem 1.

Specifically, to approach the cumulative reward R_i^t , the goal of the critic network Q_i^π is to minimize the sum of square loss $L(\varepsilon_i)$, between the critic network output $\hat{Q}_i$ and R_i^t , i.e., $L(\varepsilon_i) = \min_{\varepsilon_i} \mathbb{E}[(\hat{Q}_i - R_i)^2]$. Thus, the gradient of critic networks $\nabla_{\varepsilon_i} L(\varepsilon_i)$ is calculated by

$$\nabla_{\varepsilon_i} L(\varepsilon_i) = \mathbb{E}[(Q_i^\pi(s, a_{1,\dots,Ni}) - y) \cdot \nabla_{\varepsilon_i} Q_i^\pi(s, a_{1,\dots,Ni})] \quad (\text{S-1})$$

where $y = r_i^t + \gamma Q_i^\pi(s^{t+1}, a_{1,\dots,N}^{t+1}|_{a_i^{t+1}=\pi_i(o_i)})$, which is an approximate value calculated by the temporal difference (TD).

Theorem: For the learning process of critic networks by $\varepsilon_i^{k+1} = \varepsilon_i^k - \alpha_c \nabla_{\varepsilon_i} L(\varepsilon_i)$, where α_c is the critic networks' learning rate. Assuming that the time interval Δt is constant, the parameters ε_i^k of each defender asymptotically converges to the optimum value $\varepsilon_i^* = \underset{\varepsilon_i}{\operatorname{argmin}} E[(\hat{Q}_i - R_i)^2]$ with iteration index $k \rightarrow \infty$, if the learning rate α_c satisfies

$$0 < \alpha_c < 2. \quad (\text{S-2})$$

In other words, the error of the critical neural network can converge.

Proof: For defender i , denote the critic network output $\hat{Q}_i^k(x)$ by $\hat{Q}_i^k(x) = (\varepsilon_i^k)^T \phi(x)$, where x is the critic network input, including the states s and the actions a of all the defenders $\mathcal{V}_d$, and

$$\phi(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}} \quad (\text{S-3})$$

is the hyperbolic activation function of critic networks. For conciseness, we omit the symbol “ (x) ” in the following context. The error e_i^k between the reward prediction $\hat{Q}_i^k$ and the real reward R_i writes $e_i^k = \hat{Q}_i^k - R_i$, $L_i^k = \frac{1}{2} (e_i^k)^2$, where L_i^k is the square loss for obtaining gradients. To minimize L_i^k , ε_i^k is updated by

$$\varepsilon_i^{k+1} = \varepsilon_i^k - \alpha_c \partial L_i^k / \partial \varepsilon_i^k = \varepsilon_i^k - \alpha_c e_i^k \quad (\text{S-4})$$

where α_c is the learning rate of critic networks, and $\alpha_c \geq 0$. Then the error of the parameters $\bar{\varepsilon}_i^k$ can be expressed as $\bar{\varepsilon}_i^k = \varepsilon_i^k - \varepsilon_i^*$. Thus, one has $\bar{\varepsilon}_i^{k+1} = \bar{\varepsilon}_i^k - \alpha_c e_i^k \phi$. Now, design a discrete-time Lyapunov candidate as $V(k) = (\bar{\varepsilon}_i^k)^T \bar{\varepsilon}_i^k$, whose difference writes:

$$\begin{aligned} \Delta V(k) &= V(k+1) - V(k) \\ &= (\bar{\varepsilon}_i^{k+1})^T \bar{\varepsilon}_i^{k+1} - (\bar{\varepsilon}_i^k)^T \bar{\varepsilon}_i^k \\ &= [(\bar{\varepsilon}_i^k)^T - \alpha_c \phi^T (e_i^k)^T] (\bar{\varepsilon}_i^k - \alpha_c e_i^k \phi) - (\bar{\varepsilon}_i^k)^T \bar{\varepsilon}_i^k \\ &= -\alpha_c \phi^T (e_i^k)^T \bar{\varepsilon}_i^k - \alpha_c (\bar{\varepsilon}_i^k)^T e_i^k \phi + \alpha_c^2 (e_i^k \phi)^T e_i^k \phi \\ &= -\alpha_c \phi^T (e_i^k)^T (\varepsilon_i^k - \varepsilon_i^*) - \alpha_c (\varepsilon_i^k - \varepsilon_i^*)^T e_i^k \phi \\ &\quad + \alpha_c^2 \|e_i^k\|^2 \|\phi\|^2 \\ &= \alpha_c \|e_i^k\|^2 (-2 + \alpha_c \|\phi\|^2). \end{aligned} \quad (\text{S-5})$$

According to (S-2) and (S-3), one has $-2 + \alpha_c \|\phi\|^2 < 0$, and hence $\Delta V < 0$, implying that ε_i^k asymptotically converges to ε_i^* . The error of the critical neural network can converge to 0. ■