VC dimension and Eluder dimension

Xuhui Liu

LAMDA, NJUAI

July 22, 2020

Outline

Table of Contents

VC dimension

Definition 4. An action a is *VC-independent* of $\tilde{\mathscr{A}} \subseteq \mathscr{A}$ if for any $f, \tilde{f} \in \mathscr{F}$, there exists some $\bar{f} \in \mathscr{F}$, which agrees with f on a and with $\tilde{f}$ on $\tilde{\mathscr{A}}$; that is, $\bar{f}(a) = f(a)$ and $\bar{f}(\tilde{a}) = \tilde{f}(\tilde{a})$ for all $\tilde{a} \in \tilde{\mathscr{A}}$. Otherwise, a is *VC-dependent* on $\tilde{\mathscr{A}}$.

Definition 5. The VC dimension of a class of binary-valued functions with domain $\mathscr{A}$ is the largest cardinality of a set $\tilde{\mathscr{A}} \subseteq \mathscr{A}$ such that every $a \in \tilde{\mathscr{A}}$ is VC-independent of $\tilde{\mathscr{A}} \backslash \{a\}$.

- Only one function in $\mathscr{F}$.
  In this case, VC dimension $= 1$. But applying the new definition we have VC dimension $= \infty$.
- Only two functions in $\mathscr{F}$.
  VC dimension $= 1$ or 2. But applying the new definition we have VC dimension $= \infty$.

Theorem

If $0, 1 \in \mathscr{F}$, then the two definitions are equal.

Proof.

It's trivial that the old definition implies the new definition. We focus on the other direction.

$\forall f \in 2^{\mathscr{A}}$, where $|\mathscr{A}| = \dim VC$. Let $X = \{a : f(a) = 1\}$, by the definition, we have

$$\bigcup_{a \in X} \{g : g(Y) = 1, g(\mathscr{A} \backslash Y) = 0, Y = Y \cup \{a\}\} \subset \mathscr{F}$$

whereas $f \in \bigcup_{a \in X} \{g : g(Y) = 1, g(\mathscr{A} \backslash Y) = 0, Y = Y \cup \{a\}\} \subset \mathscr{F}$, so $f \in \mathscr{F}$. $\square$

Table of Contents

Problem formulation

- The T-proiod regret of a policy π is defined by

$$\text{Regret}(T, \pi, \theta) \sum_{t=1}^{T} \mathbb{E}[\max_{a \in \mathscr{A}_t} f_\theta(a) - f_\theta(A_t) | \theta].$$

- The T-period Bayesian regret is defined by $\mathbb{E}[\text{Regret}(T, \pi, \theta)]$, where the expectation is taken with respect to the prior distribution over θ. Hence,

$$\text{BayesRegret}(T, \pi) = \sum_{t=1}^{T} \mathbb{E}[\max_{a \in \mathscr{A}_t} f_\theta(a) - f_\theta(A_t)].$$

Assumption

- Assumption 1. For all $f \in \mathscr{F}$ and $a \in \mathscr{A}$, $f(a) \in [0, C]$.
- Assumption 2. For all $t \in \mathbb{N}$, $R_t - f\theta(A_t)$ conditioned on (H_t, θ, A_t) is σ-sub-Gaussian.

Proof sketch

Let $A_t^* = \max_{a \in \mathscr{A}_t} f_\theta(a)$, we have

$$\text{BayesRegret}(T, \pi) = \sum_{t=1}^{T} \mathbb{E}[f_\theta(A_t^*) - f_\theta(A_t)].$$

Proof sketch

Step 1: **Regret decomposition.**
If $f_\theta(A_t) \in [L_t(A_t), U_t(A_t)]$, then

$$\mathbb{E}[f_\theta(A_t^*) - f_\theta(A_t)] \leq \mathbb{E}[U_t(A_t^*) - f_\theta(A_t)] \leq \mathbb{E}[U_t(A_t) - f_\theta(A_t)].$$

Else

$$f_\theta(A_t^*) - f_\theta(A_t) \leq C.$$

$\Rightarrow \text{BayesRegret} \leq \mathbb{E}\sum_{t=1}^{T}[(U_t(A_t) - f_\theta(A_t)] + C\sum_{t=1}^{T}\mathbb{P}(f_\theta(A_t^*) > U_t(A_t^*))$.

Proof sketch

Step 2: **Build confidence set.**
Minimize MSE loss $\rightarrow f_\theta(A_t) \in [L_t(A_t), U_t(A_t)]$ with high probability.

Proof sketch

Step 3: **Bound the confidence set interval.**
We have $U_t(A_t) - f_\theta(A_t) \leq \omega_t(A_t) := U_t(A_t) - L_t(A_t)$, and bound $\sum_{t=1}^T \omega_t(A_t)$ to conclude the proof.

Step 1:

PROPOSITION 1. *For any UCB sequence* $\{U_t \mid t \in \mathbb{N}\}$,

$$\text{BayesRegret}(T, \pi^{PS}) = \mathbb{E}\sum_{t=1}^{T}[U_t(A_t) - f_\theta(A_t)] + \mathbb{E}\sum_{t=1}^{T}[f_\theta(A_t^*) - U_t(A_t^*)]$$

for all $T \in \mathbb{N}$.

$$L_{2,t}(f) = \sum_{1}^{t-1} (f(A_t) - R_t)^2$$

$$||\bar{f} - \underline{f}||_{2,E_t}^2 = \sum_{k=1}^{t-1} (\bar{f} - \underline{f})^2 (A_k)$$

Step 2:

LEMMA 3. *For any* $\delta > 0$ *and* $f \colon \mathscr{A} \mapsto \mathbb{R}$, *with probability at least* $1 - \delta$,

$$L_{2,t}(f) \geq L_{2,t}(f_\theta) + \frac{1}{2}\|f - f_\theta\|_{2,E_t}^2 - 4\sigma^2 \log(1/\delta)$$

simultaneously for all $t \in \mathbb{N}$.

$$\beta_t^*(\mathscr{F}, \delta, \alpha) := 8\sigma^2 \log(N(\mathscr{F}, \alpha, \|\cdot\|_\infty)/\delta) + 2\alpha t(8C + \sqrt{8\sigma^2 \ln(4t^2/\delta)}).$$

PROPOSITION 6. *For all* $\delta > 0$ *and* $\alpha > 0$, *if*

$$\mathscr{F}_t = \{f \in \mathscr{F} \colon \|f - \hat{f}_t^{LS}\|_{2,E_t} \leq \sqrt{\beta_t^*(\mathscr{F}, \delta, \alpha)}\}$$

for all $t \in \mathbb{N}$, *then*

$$\mathbb{P}\left(f_\theta \in \bigcap_{t=1}^{\infty} \mathscr{F}_t\right) \geq 1 - 2\delta.$$

Step 3:

LEMMA 4. *For all* $T \in \mathbb{N}$, *if* $\inf_{\rho \in \mathcal{F}_\tau} f_\rho(a) \le f_\theta(a) \le \sup_{\rho \in \mathcal{F}_\tau} f_\rho(a)$ *for all* $\tau \in \mathbb{N}$ *and* $a \in \mathcal{A}$ *with probability at least* $1 - 1/T$, *then*

$$\text{BayesRegret}(T, \pi^{PS}) \le C + \mathbb{E} \sum_{t=1}^{T} w_{\mathcal{F}_t}(A_t).$$

Step 3:

PROPOSITION 8. *If* $(\beta_t \geq 0 \mid t \in \mathbb{N})$ *is a nondecreasing sequence and* $\mathcal{F}_t := \{f \in \mathcal{F} : \|f - \hat{f}_t^{LS}\|_{2,E_t} \leq \sqrt{\beta_t}\}$, *then*

$$\sum_{t=1}^{T} \mathbf{1}(w_{\mathcal{F}_t}(A_t) > \epsilon) \leq \left(\frac{4\beta_T}{\epsilon^2} + 1\right) \dim_E(\mathcal{F}, \epsilon)$$

for all $T \in \mathbb{N}$ *and* $\epsilon > 0$.

LEMMA 5. *If* $(\beta_t \geq 0 \mid t \in \mathbb{N})$ *is a nondecreasing sequence and* $\mathcal{F}_t := \{f \in \mathcal{F} : \|f - \hat{f}_t^{LS}\|_{2,E_t} \leq \sqrt{\beta_t}\}$, *then*

$$\sum_{t=1}^{T} w_{\mathcal{F}_t}(A_t) \leq 1 + \dim_E(\mathcal{F}, T^{-1})C + 4\sqrt{\dim_E(\mathcal{F}, T^{-1})\beta_T T}$$

for all $T \in \mathbb{N}$.

Step 3:

Definition 1. The *Kolmogorov dimension* of a function class $\mathcal{F}$ is given by

$$\dim_K(\mathcal{F}) = \limsup_{\alpha \downarrow 0} \frac{\log N(\mathcal{F}, \alpha, \|\cdot\|_\infty)}{\log(1/\alpha)}.$$

In particular, we have the following result.

Proposition 7. *For any fixed class of functions* $\mathcal{F}$,

$$\beta_t^*(\mathcal{F}, 1/t^2, 1/t^2) = 16(1 + o(1) + \dim_K(\mathcal{F})) \log t.$$

Proposition 10. *For any fixed class of functions* $\mathcal{F}$,

$$\text{BayesRegret}(T, \pi^{PS}) \leq 1 + [\dim_E(\mathcal{F}, T^{-1}) + 1]C + 16\sigma\sqrt{\dim_E(\mathcal{F}, T^{-1})(1 + o(1) + \dim_K(\mathcal{F})) \log(T) T}$$

Thanks!