PAC-Bayesian Bounds on Rate-Efficient Machines

Alhabib Abbas[1], Yiannis Andreopoulos[2],

[1] University College London, alhabib.abbas.13@ucl.ac.uk
[2] University College London, i.andreopoulos@ucl.ac.uk

July 6, 2022

Motivation

When are **PAC** Classifiers Important

- Assume we have a sensor on a car.
- If the sensor observes a human, the car should swerve and compromise itself.
- If the object is not human, the car should not swerve.
- **We would like mathematical guarantees on how often the car misclassifies objects as "not human".**

When are **Rate-Efficient** Classifiers Important

- Distributed systems constrain processable data.
- Sensors compress inputs, introducing reconstruction noise.
- Classifier decisions are affected by noise.
- There is a tradeoff between latency and performance.
- **We would like mathematical guarantees on the minimum amount of information required for accurate inference.**

Rate-Efficient PAC Classifiers

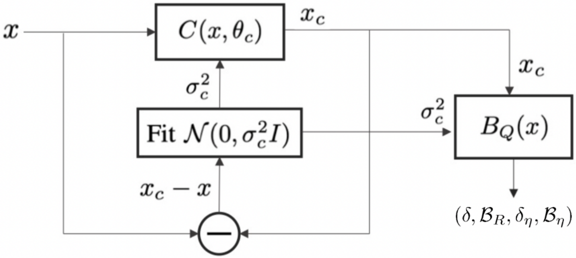

- A PAC classifier $\mathcal{M}(x_c)$ classifies x and tells us $\mathcal{M}(x) = y$, with accuracy $1 - \epsilon$, and with probability δ.
- A rate-efficient $\mathcal{M}(x_c)$ guarantees $\mathcal{M}(x) = \mathcal{M}(x_c)$, at a rate of $1 - \eta$, with probability δ_η.

PAC-Bayesian Bounds on Noise Invariance

Noise Invariance

Definition

For any source domain D, for any classification model $\mathcal{M}$, and for any noise vector $\boldsymbol{n} \sim \mathcal{N}(\mathbf{0}, \sigma_c^2 I)$ modelling perturbations on $\boldsymbol{x}$, noise invariance η^D quantifies the probability of output change due to $\boldsymbol{n}$:

$$\eta^D = \mathop{\mathbb{E}}_{x \sim D} \mathop{Pr}_{\boldsymbol{n} \sim \mathcal{N}} \Big(\mathcal{M}(\boldsymbol{x}) \neq \mathcal{M}(\boldsymbol{x} + \boldsymbol{n}) \Big) \tag{1}$$

Bounding Noise Invariance of the Majority Vote

Theorem

For any majority vote classifier defining a posterior Q over normalised linear voters $x'_i \in \mathcal{S} \subseteq \mathcal{X}$ where $h_i(\boldsymbol{x}) = y_i x'_i \boldsymbol{x}^\top$, and when $\boldsymbol{\omega} = \mathbb{E}_{x'_i \sim Q}\, \mathbb{E}_{x'_j \sim Q}\, \boldsymbol{x}'_j \boldsymbol{x}^\top \boldsymbol{x}'_i$, invariance coefficients η_Q^D are simplified to:

$$\eta_Q^D = \mathop{\mathbb{E}}_{\boldsymbol{x} \sim D} \frac{1}{2}\left[1 + \text{erf}\left(\frac{a_Q(x)}{\sqrt{\boldsymbol{\omega}\sigma_c^2 I \boldsymbol{\omega}^\top}\sqrt{2}}\right)\right] \tag{2}$$

Bounding Noise Invariance of the Majority Vote

Theorem

For any source distribution D, for any prior P^2 on the hypothesis set $\mathcal{H}^2$, for any posterior Q^2 learned by observing $\mathcal{S} \sim D^m$, and for any arbitrary probability $\delta_\eta \in (0, 1]$:

$$\Pr_{\mathcal{S} \sim D} \left(\mathrm{kl}(\eta_Q^{\mathcal{S}} || \eta_Q^{D}) \leq \frac{1}{m} \left[2KL(Q||P) + \ln \frac{\xi(m)}{\delta_\eta} \right] \right) \geq 1 - \delta_\eta \quad (3)$$

Concluding Results

Concluding Results

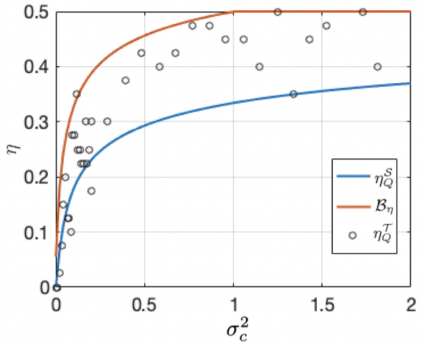

- Bounds $\mathcal{B}_\eta$ are reliable across all values of σ_c^2.
- Symmetric noise sources necessarily saturate $\mathcal{B}_\eta$ and η at 0.5.
- Values of σ_c^2 correlate inversely with x_c bitrates.

Thank You