Information at the Interface between Machine Learning and Game Theory

Nicolò Cesa-Bianchi

Università degli Studi di Milano

Algorithmic complexity and universal prediction

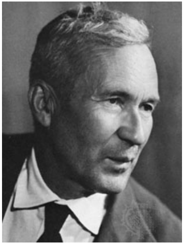


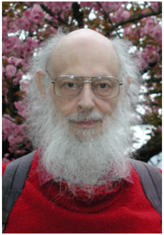

- Kolmogorov complexity K
  Information content of a binary sequence or integer

- Universal measure M
  Mixture over all continuous enumerable measures $\mu_1, \mu_2, \ldots$

$$M = \sum_i 2^{-K(i)} \mu_i$$

The predictive power of M

- Any continuous measure μ corresponds to a sequential probabilistic predictor for binary sequences $\mathbf{y} = (y_1, y_2, \dots)$
- How quickly does $M\big(y_t \mid y_1, \dots, y_{t-1}\big) = M(y_t \mid \mathbf{y}_{t-1})$ converge to 1?

$$\begin{aligned}\sum_t \big(1 - M(y_t \mid \mathbf{y}_{t-1})\big)^2 &\leqslant -\frac{1}{2} \sum_t \ln M(y_t \mid \mathbf{y}_{t-1}) \\ &= -\frac{1}{2} \ln M(\mathbf{y}) \\ &\simeq K(\mathbf{y})\end{aligned}$$

- Solomonoff's classical result: M can predict a random sequence drawn from any computable measure μ_i almost as well as μ_i itself

Theory of repeated games

Play a zero-sum game repeatedly against a possibly suboptimal opponent

Prediction with expert advice

- Experts are given classes of probabilistic predictors
  for example, probabilistic FSA with a fixed number of states
- Devise a strategy that predicts any sequence nearly as well as the best expert for that sequence
- Efficiency, nonasymptotic bounds, optimal rates

Example of regret bounds

Fix a class $\mathcal{F} = \{\mu_1, \mu_2, \dots\}$ of probabilistic predictors (experts) together with some prior μ_0 on $\mathcal{F}$

There exists a probabilistic prediction strategy whose expected number of mistakes $\widehat{L}_T$ on any binary sequence $\boldsymbol{y} = (y_1, \dots, y_T)$ satisfies

$$\underbrace{\widehat{L}_T - L_T(i)}_{\text{REGRET}} \simeq \sqrt{T \ln \frac{1}{\mu_0(i)}} \qquad i = 1, 2, \dots$$

where $L_T(i)$ is expected number of mistakes of μ_i on $\boldsymbol{y}$

Historical background

- Sequential adaptive compression: predicting as the best finite-state automata under logarithmic loss (Lempel and Ziv, 1976)
- Gambling and portfolio selection: linear experts in the simplex (Cover, 1965, 1974 and 1991)
- FSA prediction: improved analysis for finite-state automata (Feder, Merhav and Gutman, 1992)
- Linear classification: experts as functions in Euclidean, Hilbert and Banach spaces (Novikov, 1962; Freund and Schapire, 1999; Vovk, 2007)

On-line linear classification

- Sequential classification of $(\boldsymbol{x}_1, y_1), (\boldsymbol{x}_2, y_2), \dots \in \mathbb{R}^d \times \{-1, +1\}$
- Linear classifiers, $\text{SGN}\big(\boldsymbol{w}^\top \boldsymbol{x}_t\big)$

On-line linear classification

- Sequential classification of $(x_1, y_1), (x_2, y_2), \dots \in \mathbb{R}^d \times \{-1, +1\}$
- Linear classifiers, $\mathrm{SGN}(\boldsymbol{w}^\top \boldsymbol{x}_t)$

- There are simple reductions from online linear classification to prediction with expert advice
- This yields regret bounds for popular algorithms such as Perceptron (Rosenblatt, 1958) and Winnow (Littlestone, 1989)
- Different approaches for classification in normed spaces

Some issues

- Vovk re-derived basic measure-theoretic phenomena (large numbers, central limit, iterated logs) in a game-theoric framework
- Extend this program to statistical learning/inference

Some issues

- Vovk re-derived basic measure-theoretic phenomena (large numbers, central limit, iterated logs) in a game-theoric framework
- Extend this program to statistical learning/inference
- Transfer from game theory to machine learning: concepts (e.g., partial feedback) and analysis (e.g., regret bounds)

Some issues

- Vovk re-derived basic measure-theoretic phenomena (large numbers, central limit, iterated logs) in a game-theoric framework
- Extend this program to statistical learning/inference
- Transfer from game theory to machine learning: concepts (e.g., partial feedback) and analysis (e.g., regret bounds)
- Regret analysis in classification/regression: what are the limits?