PAC-Bayesian Inequalities for Martingales

Yevgeny Seldin[1,4] **François Laviolette**[2] **Nicolò Cesa-Bianchi**[3] **John Shawe-Taylor**[4] **Peter Auer**[5]

[1]Max Planck Institute for Intelligent Systems, Tübingen, Germany

[2]Département d'informatique, Université Laval, Québec, Canada

[3]Dipartimento di Informatica, Università degli Studi di Milano, Italy

[4]Department of Computer Science, University College London, UK

[5]Chair for Information Technology, Montanuniversität Leoben, Austria

seldin@tuebingen.mpg.de, francois.laviolette@ift.ulaval.ca, nicolo.cesa-bianchi@unimi.it, jst@cs.ucl.ac.uk, auer@unileoben.ac.at

Abstract

We present a set of high-probability inequalities that control the concentration of weighted averages of multiple (possibly uncountably many) simultaneously evolving and interdependent martingales. Our results extend the PAC-Bayesian analysis in learning theory from the i.i.d. setting to martingales opening the way for its application in reinforcement learning and other interactive learning domains, as well as many other domains in probability theory and statistics, where martingales are encountered.

We also present a comparison inequality that bounds the expectation of a convex function of a martingale difference sequence shifted to the $[0,1]$ interval by the expectation of the same function of independent Bernoulli variables. This inequality is applied to derive a tighter analog of Hoeffding-Azuma's inequality.

For the complete paper see Seldin et al. (2012).

$$\mathcal{H}\left\{\begin{array}{ccccccc}
\bar M_1(1) & \underset{\searrow}{\rightleftarrows} & \bar M_2(1) & \underset{\searrow}{\rightleftarrows} & \cdots & \underset{\searrow}{\rightleftarrows} & \bar M_n(1) \\
\updownarrow & & \updownarrow & & & & \updownarrow \\
\bar M_1(2) & \overset{\nearrow}{\underset{\searrow}{\rightleftarrows}} & \bar M_2(2) & \overset{\nearrow}{\underset{\searrow}{\rightleftarrows}} & \cdots & \overset{\nearrow}{\underset{\searrow}{\rightleftarrows}} & \bar M_n(2) \\
\updownarrow & & \updownarrow & & & & \updownarrow \\
\vdots & & \vdots & & \ddots & & \vdots \\
\updownarrow & & \updownarrow & & & & \updownarrow \\
\bar M_1(h) & \overset{\nearrow}{\underset{\searrow}{\rightleftarrows}} & \bar M_2(h) & \overset{\nearrow}{\underset{\searrow}{\rightleftarrows}} & \cdots & \overset{\nearrow}{\underset{\searrow}{\rightleftarrows}} & \bar M_n(h) \\
\updownarrow & & \updownarrow & & & & \updownarrow \\
\vdots & & \vdots & & \ddots & & \vdots
\end{array}\right.$$

$$\underset{\textit{time}}{\xrightarrow{\hspace{3em}}}$$

Figure 1: Illustration of an infinite set of simultaneously evolving and interdependent martingales. For a fixed h, the sequence $\bar M_1(h), \bar M_2(h), \ldots, \bar M_n(h)$ is a martingale. $\mathcal{H}$ is a set (possibly uncountably infinite) that indexes the individual martingales. The arrows represent the dependencies between the values of the martingales: the value of the h-th martingale at time i, denoted by $\bar M_i(h)$, depends on $\bar M_j(h')$ for all $j \le i$ and $h' \in \mathcal{H}$ (everything that is "before" and "concurrent" with $\bar M_i(h)$ in time; some of the arrows are omitted for clarity). A mean value of the martingales with respect to a distribution ρ over $\mathcal{H}$ (or a probability density function, if $\mathcal{H}$ is uncountably infinite) is given by $\langle \bar M_n, \rho\rangle$. Our high-probability inequalities bound $|\langle \bar M_n, \rho\rangle|$ simultaneously for a large class of ρ.

References

Yevgeny Seldin, François Laviolette, Nicolò Cesa-Bianchi, John Shawe-Taylor, and Peter Auer. PAC-Bayesian inequalities for martingales. *IEEE Transactions on Information Theory*, 2012. Accepted. Preprint available at http://arxiv.org/abs/1110.6886.