36-771 Fall 2018

Martingales 1 : Concentration inequalities

Stitching for subGamma/subGaussian boundaries

Lecturer : Aaditya Ramdas

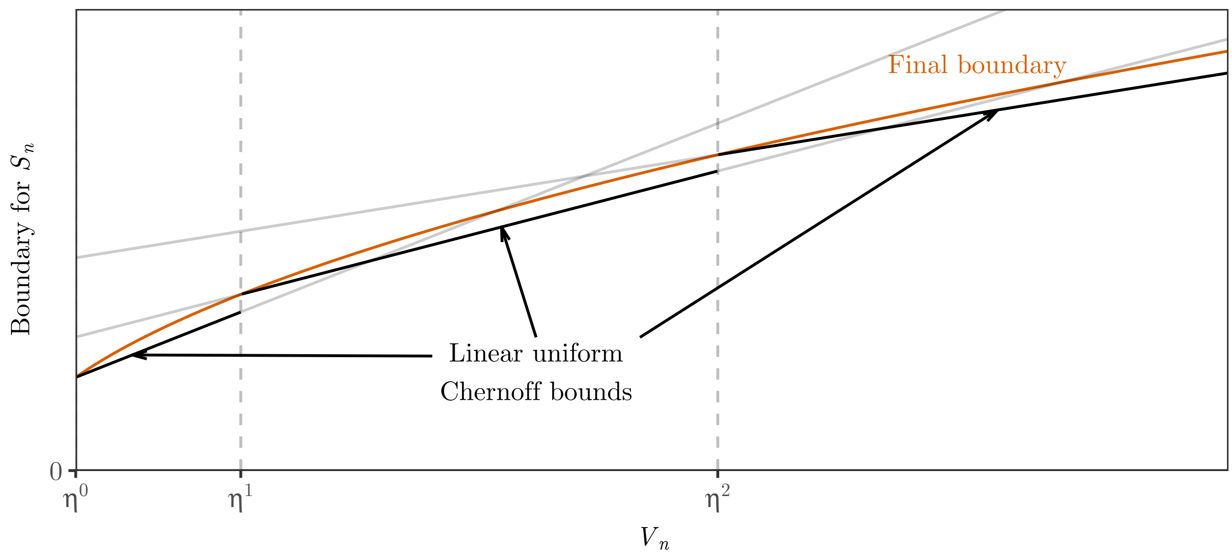

Figure 1: Illustration of stitching together linear boundaries to construct a curved boundary. We break time into geometrically-spaced epochs $\eta^k \leq V_t < e^{k+1}$, construct a linear uniform bound using Lemma 1 (previous lecture) optimized for each epoch, and take a union bound over all crossing events. The final boundary is a smooth analytical upper bound to the piecewise linear bound.

1 Analytical bounds: the stitching method

The idea behind Theorem 1is to divide intrinsic time into geometrically spaced epochs, $\eta^k \leq V_t < \eta^{k+1}$ for some $\eta > 1$. We construct a linear boundary within each epoch using Lemma 1 (previous lecture) and take a union bound over crossing events of the different boundaries. The resulting, piecewise-linear boundary may then be upper bounded by a smooth, concave function. Figure 1 illustrates the construction.

The boundary shape is determined by choosing the function h and setting the nominal crossing probability in the k^{th} epoch to equal $\alpha/h(k)$. Then Theorem 1gives a curved boundary which grows at a rate $\mathcal{O}(\sqrt{V_t \log h(\log_\eta V_t)})$ as $V_t \uparrow \infty$. The more slowly $h(k)$ grows as $k \uparrow \infty$, the more slowly the resulting boundary will grow as $V_t \uparrow \infty$. A simple choice is exponential growth, $h(k) = \eta^{sk}/(1 - \eta^{-s})$ for some $s > 1$, yielding $\mathcal{S}_\alpha(v) = \mathcal{O}(\sqrt{v \log v})$.

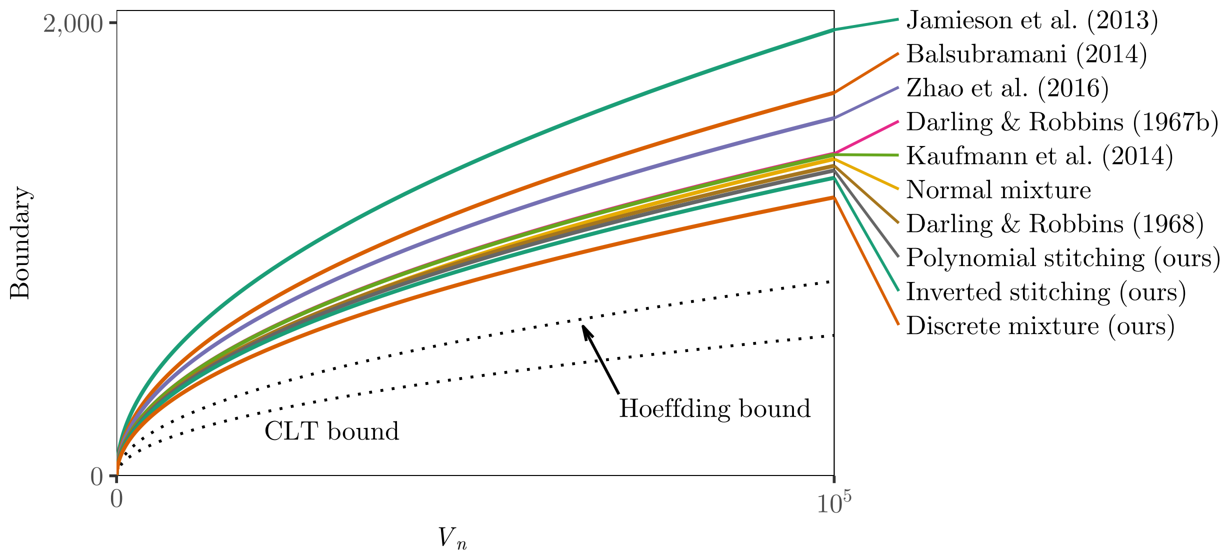

Figure 2: Finite LIL bounds for independent 1-sub-Gaussian observations, $\alpha = 0.025$. The dotted lines show fixed-sample Hoeffding bound $\sqrt{2t \log \alpha^{-1}}$, which is nonasymptotically pointwise valid but not uniformly valid, and the fixed-sample CLT bound $z_{1-\alpha}\sqrt{t}$ which is asymptotically pointwise valid. Polynomial stitching uses Theorem 1with $\eta = 2.04$ and $h(k) = (k+1)^{1.4}\zeta(1.4)$. The inverted stitching boundary is $1.7\sqrt{V_t(\log(1 + \log V_t) + 3.5)}$, using the inverted stitching theorem (later class) with $\eta = 2.99$, $v_{\max} = 10^{20}$, and error rate 0.815α to account for finite horizon. Discrete mixture uses the discrete mixture theorem (later class) with $f(\lambda) \propto 1/\lambda \log^{1.4}(1/\lambda), \eta = 1.1$, and $\lambda_{\max} = 4$. The normal mixture bound (later class) uses $\rho = 0.129$. See Howard et al. (2018*b*) for details.

1.1 Polynomial stitching and finite LIL bounds

Recall that we used $\zeta(s) = \sum_{k=1}^{\infty} s^{-k}$ to represent the Riemann zeta function. Choosing $h(k) = (k+1)^s \zeta(s)$ for some $s > 1$ in Theorem 1yields $\mathcal{S}_\alpha(v) \sim \sqrt{(2+\delta)v \log\log v})$, where we may attain any $\delta > 0$ by taking η and s sufficiently close to one, coming arbitrarily close to the lower bound furnished by the classical LIL. Uniform bounds achieving this iterated logarithm growth rate are known as *finite LIL bounds*. One may substitute a series converging yet more slowly; for example, $h(k) \propto (k+2) \log^{\mathfrak{s}}(k+2)$ for $s > 1$ yields

$$\log h(\log_\eta V_t) = \log\log_\eta(\eta^2 V_t) + s \log\log\log_\eta(\eta^2 V_t) + \log\left(\frac{\log^{1-s}(3/2)}{s-1}\right), \qquad (1)$$

matching related analysis in Darling & Robbins (1967), Robbins & Siegmund (1969), Robbins (1970), and Balsubramani (2014). In practice, the bound (1) appears to behave like bound (**??**) with worse constants. However, the fact that the stitching approach can recover key theoretical results like these gives some indication of its power. Figure 2 compares our polynomial stitching bound for 1-sub-Gaussian increments to a variety of bounds from the literature; our bound shows a slight improvement. We also include a numerically-computed discrete mixture bound with a mixture distribution roughly corresponding to $h(k) \propto (k+1)^{1.4}$, as described later. This acts as a lower bound and shows that not too much is lost by the approximations involved in the stitching construction.

1.2 Why do we get tighter finite LIL bounds than past work?

The idea of taking a union bound over geometrically spaced epochs is standard in the proof of the classical law of the iterated logarithm (Durrett 2017, Theorem 8.5.1). The idea has been extended to finite-time bounds by Darling & Robbins (1967), Jamieson et al. (2014), Kaufmann et al. (2014), and Zhao et al. (2016), usually when the observations are independent and sub-Gaussian. Of course, Theorem 1generalizes these constructions much beyond the independent sub-Gaussian case, but it also achieves tighter constants for the sub-Gaussian setting. Here, we briefly discuss how the improved constants arise.

Both Jamieson et al. (2014) and Zhao et al. (2016) construct a constant boundary rather than a linear increasing boundary over each epoch. They apply Doob's maximal inequality for submartingales (Durrett 2017, Theorem 4.4.2), as in Hoeffding (1963, eq. 2.17), to obtain boundaries similar to that of Freedman (1975). As illustrated in Howard et al. (2018*a*, Figure 2), the linear bounds from Lemma 1 (previous lecture) are stronger than corresponding Freedman-style bounds, and the additional flexibility yields tighter constants.

Both Darling & Robbins (1967) and Kaufmann et al. (2014) use linear boundaries within each epoch analogous to those of Lemma 1 (previous lecture). Both methods share a great deal in common with ours, and Darling & Robbins give consideration to general cumulant-generating functions. Recall from Lemma 1 (previous lecture) that such linear boundaries may be chosen to optimize for some fixed time $V_t = m$. Our method chooses the linear boundary within each epoch to be optimal at the geometric center of the epoch, i.e., at $V_t = \eta^{k+1/2}$, so that at both epoch endpoints the boundary will be equally "loose", that is, equal multiples of $\sqrt{V_t}$. Darling & Robbins choose the boundaries to be tangent at the start of the epoch, hence their boundary is looser than ours at the end of the epoch. Kaufmann et al. choose the boundary as we do, but appear to incur more looseness in the subsequent inequalities used to construct a smooth upper bound.

References

Balsubramani, A. (2014), 'Sharp Finite-Time Iterated-Logarithm Martingale Concentration', *arXiv:1405.2639 [cs, math, stat]* .

Darling, D. A. & Robbins, H. (1967), 'Iterated Logarithm Inequalities', *Proceedings of the National Academy of Sciences* **57**(5), 1188–1192.

Durrett, R. (2017), *Probability: Theory and Examples*, 5a edn.

Freedman, D. A. (1975), 'On Tail Probabilities for Martingales', *The Annals of Probability* **3**(1), 100–118.

Hoeffding, W. (1963), 'Probability Inequalities for Sums of Bounded Random Variables', *Journal of the American Statistical Association* **58**(301), 13–30.

Howard, S. R., Ramdas, A., McAuliffe, J. & Sekhon, J. (2018*a*), 'Exponential line-crossing inequalities', *arXiv:1808.03204 [math]* .

Howard, S. R., Ramdas, A., McAuliffe, J. & Sekhon, J. (2018*b*), 'Uniform, nonparametric, non-asymptotic confidence sequences', *arXiv:1808.08240 [math]* .

Jamieson, K., Malloy, M., Nowak, R. & Bubeck, S. (2014), lil' UCB: An Optimal Exploration Algorithm for Multi-Armed Bandits, *in* 'Proceedings of The 27th Conference on Learning Theory', Vol. 35 of *Proceedings of Machine Learning Research*, pp. 423–439.

Kaufmann, E., Cappé, O. & Garivier, A. (2014), 'On the Complexity of Best Arm Identification in Multi-Armed Bandit Models', *arXiv:1407.4443 [cs, stat]* .

Robbins, H. (1970), 'Statistical Methods Related to the Law of the Iterated Logarithm', *The Annals of Mathematical Statistics* **41**(5), 1397–1409.

Robbins, H. & Siegmund, D. (1969), 'Probability Distributions Related to the Law of the Iterated Logarithm', *Proceedings of the National Academy of Sciences* **62**(1), 11–13.

Zhao, S., Zhou, E., Sabharwal, A. & Ermon, S. (2016), Adaptive Concentration Inequalities for Sequential Decision Problems, *in* '30th Conference on Neural Information Processing Systems (NIPS 2016)', Barcelona, Spain.