

Comment: Monitoring Networked Applications With Incremental Quantile Estimation

Bin Yu

First of all, we would like to thank the authors for their timely paper focusing on the important area of streaming data. Recently, much attention has been paid by the statistics community to the high dimensionality or massiveness of data in the information technology age. However, streaming data represent the other important feature of the IT age, the high rate of data. Both high dimensionality and high rate require fast computation, but the real-time constraint on streaming data forces its computation to be a magnitude faster than that of the off-line or batch mode of massive data. As a result, in the absence of supercomputers, the algorithms for streaming data have to be very simple to be effective.

Chambers et al. deal with streaming data for computer system monitoring. Streaming data arise also in many other fields of science and engineering, such as astronomy, geoscience and sensor networks. Chambers et al. devise a simple and practical algorithm for updating quantiles to be used to monitor the reliability of a large system based on streamed data. Stationarity is implicitly assumed since one could argue that a good computer system should be more or less stable over time until the system is updated.

A desirable add-on to the estimated quantile of Chambers et al. is a measure of uncertainty which in the i.i.d. case is trivial because of the relationship between the variance and mean of a binomial random variable. However, it is hard to imagine that a computer system follows an i.i.d. process. The real-time constraint could make the pursuit of an uncertainty measure harder than the quantile estimation itself.

For a natural environment to be monitored by a sensor network, the variable of interest (say, temperature) is most likely to be changing over time and hence nonstationary. Fortunately, there is an easy extension of the Chambers et al. algorithm to the nonstationary case.

Bin Yu is Professor, Department of Statistics, University of California, Berkeley, Berkeley, California 94720, USA (e-mail: binyu@stat.berkeley.edu)

Because we can build the CDF and therefore the quantiles based on a moving window of data, it is applicable to nonstationary data streams. However, in this case, the data have to be kept over a duration of the size of the moving window W , in addition to the current estimate of the CDF.

Formally, let W denote the size of the moving time window which is application-specific to guarantee some stationarity of the variable within the window. Let O denote the initial block of (old) data to be removed when new data come in, K the data block kept and N the new block to be taken into account: $|W| = |O| + |K|$ and $|O| = |N|$.

Since the current empirical count of observations less than any x is a summation of the indicator function of the interval $(-\infty, x]$ over the current block of data (over K and N), it can be obtained by using the last empirical count and the summation over the old block:

$$\begin{aligned} & \sum_{t \in \text{current block}} I_{\{X_t \leq x\}} \\ &= \sum_{t \in K} I_{\{X_t \leq x\}} + \sum_{t \in N} I_{\{X_t \leq x\}} \\ &= \sum_{t \in K} I_{\{X_t \leq x\}} + \sum_{t \in N} I_{\{X_t \leq x\}} \\ & \quad + \sum_{t \in O} I_{\{X_t \leq x\}} - \sum_{t \in O} I_{\{X_t \leq x\}} \\ &= \sum_{t \in O} I_{\{X_t \leq x\}} + \sum_{t \in K} I_{\{X_t \leq x\}} \\ & \quad + \sum_{t \in N} I_{\{X_t \leq x\}} - \sum_{t \in O} I_{\{X_t \leq x\}} \\ &= \sum_{t \in \text{previous block}} I_{\{X_t \leq x\}} \\ & \quad + \sum_{t \in N} I_{\{X_t \leq x\}} - \sum_{t \in O} I_{\{X_t \leq x\}}. \end{aligned}$$

With proper scaling and weighting, the empirical CDF for the current block can be easily updated based on the