Non-parametric confidence-based cost estimation

Roberto Rossi
University of Edinburgh

February 28, 2016

Consider a stochastic constraint optimisation problem $\mathcal{P}$ [3], which without loss of generality in what follows will be formulated as a profit maximisation problem. Let $\mathcal{P}_{\text{lb}}$ and $\mathcal{P}_{\text{ub}}$ be two sampled stochastic constraint optimisation problems [2] obtained from $\mathcal{P}$ such that the optimal solution to $\mathcal{P}_{\text{lb}}$ underestimates the optimal solution to $\mathcal{P}$ with probability α and the optimal solution to $\mathcal{P}_{\text{ub}}$ overestimates the optimal solution to $\mathcal{P}$ with probability α. These sampled stochastic constraint optimisation problems can be obtained by using the notion of (α, ϑ)-solution [2].

Let ω_{lb}^i for $i = 1, \ldots, M$ be the finite-time discrete stochastic process representing the objective values obtained by repeatedly solving $\mathcal{P}_{\text{lb}}$ M times.

Let ω_{ub}^i for $i = 1, \ldots, M$ be the finite-time discrete stochastic process representing the objective values obtained by repeatedly solving $\mathcal{P}_{\text{ub}}$ M times.

Although we do not know the exact distribution of ω_{lb}^i and ω_{ub}^i, we know that these stochastic processes are stationary. In addition we know that ω_{lb}^i will underestimate the optimal solution of $\mathcal{P}$ with probability α and that ω_{lb}^i will overestimate the optimal solution of $\mathcal{P}$ with probability α.

We run the stochastic process ω_{lb}^i for $i = 1, \ldots, M$ and store the optimal profit obtained for each of these instances into an array K_{lb} sorted in ascending order; we also run the stochastic process ω_{ub}^i for $i = 1, \ldots, M$ and store the optimal profit obtained for each of these instances into an array K_{ub} sorted in ascending order.

Let $\text{bin}^{-1}(M, \alpha)$ be the inverse cumulative distribution of a binomial distribution with M trials and a success probability α; let k_{lb} be the $(1-\alpha)/2$-quantile of this distribution; finally, let k_{ub} be the $1-(1-\alpha)/2$-quantile of $\text{bin}^{-1}(M, 1-\alpha)$. With confidence α element at position k_{lb} of K_1 is a lower bound and element at position $k_{\text{ub}} + 1$ of K_{ub} is an upper bound to the true optimal cost.[1]

[1] Elements of K_i are indexed as follows: $1, \ldots, |K_i|$. Note that in statistics the k^{th}-smallest value of a statistical sample is known as k^{th} order statistic [1].

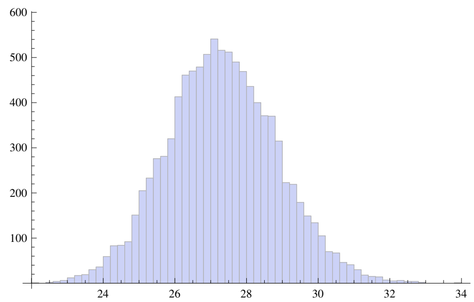

Figure 1: Empirical distribution of the k_{lb} order statistics of K_{lb}

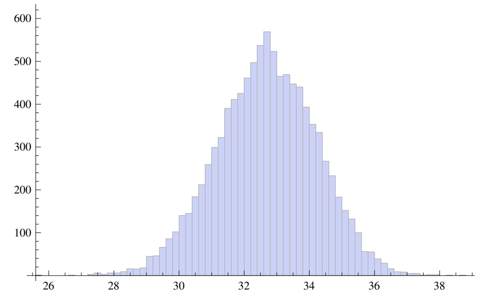

Figure 2: Empirical distribution of the k_{ub} order statistics of K_{ub}

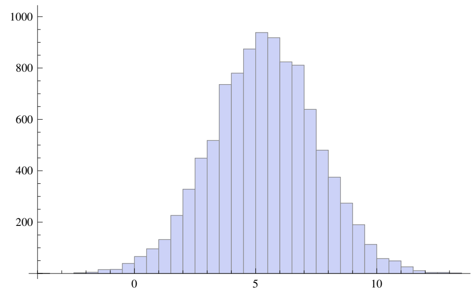

Figure 3: Empirical distribution of the difference between the k_{ub} order statistics of K_{ub} and the k_{lb} order statistics of K_{lb} for $M = 20$

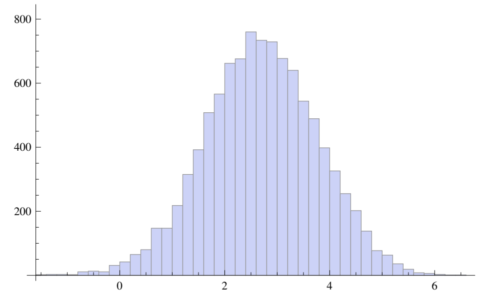

Figure 4: Empirical distribution of the difference between the k_{ub} order statistics of K_{ub} and the k_{lb} order statistics of K_{lb} for $M = 100$

Example

Assume that the value of the optimal solution to $\mathcal{P}$ is $\mu = 30$; $\sigma = 5$; G^{-1} denotes the inverse cumulative distribution function of a standard normally distributed random variable; ω^i_{lb} is normally distributed with mean $\mu_{\text{lb}} = \mu + \sigma G^{-1}(1-\alpha)$; ω^i_{ub} is normally distributed with mean $\mu_{\text{lb}} = \mu + \sigma G^{-1}(\alpha)$; $M = 20$. If we fix $M = 20$ and $\alpha = 0.9$, it follows that $k_{\text{lb}} = 16$ and $k_{\text{ub}} = 4$. Therefore element 16 of K_{lb} is a lower bound for μ and element 5 of K_{ub} is an upper bound for μ with probability α. We replicated the process 10000 times and obtained the distributions shown in Fig. 1 and Fig. 2 for the k_{lb} order statistics of K_{lb} and the k_{ub} order statistics of K_{ub}, respectively. The confidence interval obtained for μ, defined by the lower and the upper bound obtained in each run as illustrated, covers the true value of μ (i.e. 30) with frequency $0.9154 \geq \alpha$. In Fig. 3 and Fig. 4 we demonstrate how the distribution of the optimality gap varies when M takes value 20 or 100.

References

[1] Herbert A. David and H. N. Nagaraja. *Order Statistics*. Wiley-Interscience, 3 edition, August 2003.

[2] R. Rossi, B. Hnich, S. A. Tarim, and S. Prestwich. Finding (α, ϑ)-solutions via sampled SCSP. In *22nd International Joint Conference on Artificial Intelligence*, 2011.

[3] S. Tarim, Suresh Manandhar, and Toby Walsh. Stochastic constraint programming: A Scenario-Based approach. *Constraints*, 11(1):53–80–80, January 2006.

Appendix: Mathematica code

```
kLBArray={}; kUBArray={};
M=20; \[Mu]=30; \[Sigma]=5; \[Alpha]=0.9;
\[Mu]LB=InverseCDF[NormalDistribution[\[Mu],\[Sigma]],(1-\[Alpha])];
\[Mu]UB=InverseCDF[NormalDistribution[\[Mu],\[Sigma]],\[Alpha]];
counter=0; R=10000;
For[x=1,x<=R,x++,
dLB=NormalDistribution[\[Mu]LB,\[Sigma]];
dUB=NormalDistribution[\[Mu]UB,\[Sigma]];
sLB=RandomReal[dLB,M];
sUB=RandomReal[dUB,M];
sLBSorted=Sort[sLB];
sUBSorted=Sort[sUB];
lb=InverseCDF[BinomialDistribution[M,0.9],(1-\[Alpha])/2];
ub=InverseCDF[BinomialDistribution[M,0.1],1-(1-\[Alpha])/2];
kLB=sLBSorted[[lb]];
kUB=sUBSorted[[ub+1]];
kLBArray=Append[kLBArray,kLB];
kUBArray=Append[kUBArray,kUB];
If[kLB<=\[Mu] && kUB>=\[Mu],counter++];
];
N[counter/R]
Histogram[kLBArray]
Histogram[kUBArray]
Histogram[kUBArray-kLBArray]
Mean[kUBArray-kLBArray]
```