1114-60-195 **Mark L Huber*** (autotomic@gmail.com), 850 Columbia Avenue, Claremont, CA 91711.

Approximation algorithms for the normalizing constant of Gibbs distributions.

Consider a family of distributions $\{\pi_\beta\}$ where $X \sim \pi_\beta$ means that $\mathbf{P}(X = x) = \exp(-\beta H(x))/Z(\beta)$. Here $Z(\beta)$ is the normalizing constant for the density. Then $\{\pi_\beta\}$ is known as a Gibbs distribution, and $Z(\beta)$ is the partition function. This work presents new method called the paired product estimator (PPE) for approximating $Z(\beta)$. The PPE approximates the function to a specified level of relative accuracy using a number of samples that grows as $O(\ln(Z(\beta)) \ln(\ln(Z(\beta))))$ when $Z(0) \geq 1$. This is a sharp improvement for previous, similar approaches that used a much more complicated algorithm, yet used $O(\ln(Z(\beta)) \ln(\ln(Z(\beta)))^5)$ samples. (Received August 26, 2015)