

Supplementary Material:

Bucket Renormalization for Approximate Inference

A. Details for Implementation of Algorithms

Here, we provide more specific details of the algorithms implemented for the experiments. WMBE used uniform Hölder weights and reparameterization updates were run at most 100 iterations until convergence. Next, MF, BP and GBP were run for 1000 at maximum. BP and GBP additionally used damping with 0.1 ratio in order to help convergence. Finally, region graph of GBP was generated via greedily clustering neighboring factors, i.e., factors sharing an adjacent variable, into a single region until no factors were able to be clustered without violating the memory constraint given by induced-width bound *ibound*.

B. Details for UAI datasets

Here, we provide specific details of the UAI datasets used for experiments. Specifically, there exist 35 instances of Promedus GMs in with the average of 544.75 binary variables and 305.85 non-singleton hyper-edges with averaged maximum size $\max_{\alpha \in \mathcal{E}} |\alpha| = 3$. In case of Linkage GMs, there exist 17 instances with the average of 949.94 variables with averaged maximum cardinality $\max_{i \in \mathcal{V}} |\mathcal{X}_i| = 4.95$ and 727.35 non-singleton hyper-edges with averaged maximum size $\max_{\alpha \in \mathcal{E}} |\alpha| = 4.47$.

C. Comparison on Choice of Elimination Order

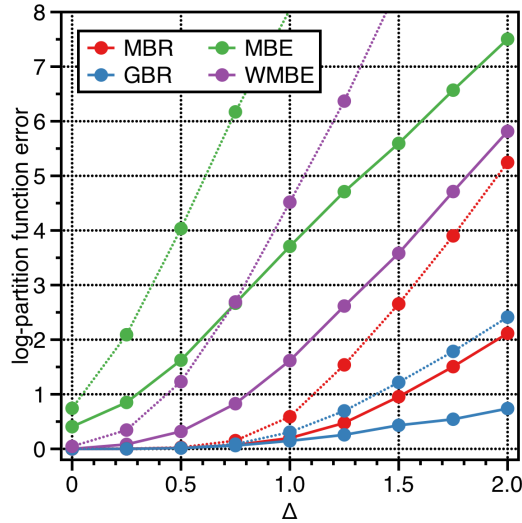


Figure 1. Performance comparisons between choosing elimination strategy via min-fill heuristics (solid) and random selection (dashed). Experiments were conducted based on the same setting as in Figure 3b.