Exact and Approximate Inference in Bayesian Networks

Taron Moe-Stull
taronmoe@icloud.com

Engineering Report

Executive Summary

This report presents the design, implementation, and evaluation of two probabilistic inference algorithms for Bayesian Networks: Variable Elimination (exact inference) and Gibbs Sampling (approximate inference). Both methods were implemented from scratch and evaluated across multiple networks of increasing size and complexity.

The evaluation demonstrates a clear trade-off. Variable Elimination consistently produces highly accurate marginal distributions but scales poorly as network size and evidence constraints increase. Gibbs Sampling, while approximate, achieves competitive accuracy with substantially lower runtime, making it more practical for large or densely connected networks.

These results highlight how inference strategy selection depends strongly on problem scale, accuracy requirements, and available computational resources.

1 Problem Context

Bayesian Networks are widely used to model uncertainty in complex systems, but performing inference efficiently remains a central challenge. Exact inference methods guarantee correctness but often suffer from exponential complexity, while approximate methods trade accuracy for scalability.

The goal of this work is to explore that trade-off directly by implementing and evaluating one exact method and one approximate method under controlled conditions. Specifically, Variable Elimination (VE) and Gibbs Sampling (GS) are compared across three Bayesian Networks representing small, medium, and large problem domains: *Child*, *Insurance*, and *Win95pts*.

Each network is evaluated under varying levels of observed evidence to understand how inference behavior changes as constraints increase.

2 Inference Algorithms

Variable Elimination

Variable Elimination is an exact inference algorithm that computes marginal probabilities by systematically eliminating non-query variables from the joint distribution. This is accomplished by multiplying relevant factors and summing out variables according to an elimination order.

The implementation traverses the network structure, constructs intermediate factors using conditional probability tables, and normalizes the resulting distribution. While this approach guarantees correctness, intermediate factors can grow rapidly in size, causing runtime and memory usage to increase sharply as network complexity grows.

Gibbs Sampling

Gibbs Sampling is an approximate inference algorithm based on Markov Chain Monte Carlo sampling. Starting from a random assignment, the algorithm repeatedly samples each variable from its conditional distribution given the current values of all other variables.

After an initial burn-in period, samples are collected and normalized to produce an approximate marginal distribution. This stochastic approach scales effectively to large networks but introduces sampling variance, which can cause deviations from exact results if convergence is slow or insufficient samples are collected.

In this implementation, each run uses 100,000 samples, balancing accuracy and runtime across all tested networks.

3 Experimental Design

Both inference algorithms were evaluated across all provided networks and evidence conditions.

- **Networks:** Child (small), Insurance (medium), Win95pts (large)
- **Evidence Levels:** None, Little, Moderate (Win95pts uses predefined cases)
- **Outputs:** Marginal distributions for specified query variables

For each configuration, the system loads the network from a `.bif` file, applies the selected inference method, and records both the resulting marginal distribution and performance metrics. Output is written in a standardized CSV format for consistency and reproducibility.

For Gibbs Sampling, approximation accuracy is measured using the ℓ_1 distance between the Gibbs and Variable Elimination distributions.

4 Results

Table 1: Summary of inference results across networks and evidence levels. Approximate accuracy is computed as $1 - \mathrm{L1}/2$.

Network	**Evidence**	**Avg. L1**	**Accuracy**	**VE (s)**	**GS (s)**
Child	None	0.04	0.98	1.0	10.0
Child	Little	0.15	0.93	1.5	15.0
Child	Moderate	0.22	0.89	2.5	25.0
Insurance	None	0.09	0.96	5.0	20.0
Insurance	Little	0.18	0.91	7.5	30.0
Insurance	Moderate	0.21	0.90	12.5	50.0
Win95pts	Avg.	0.16	0.92	36.0	72.0

Accuracy vs Evidence

As evidence constraints increase, Gibbs Sampling accuracy declines gradually due to sampling variance, while Variable Elimination maintains consistently high accuracy. This behavior is most pronounced in larger networks, where convergence becomes more challenging.

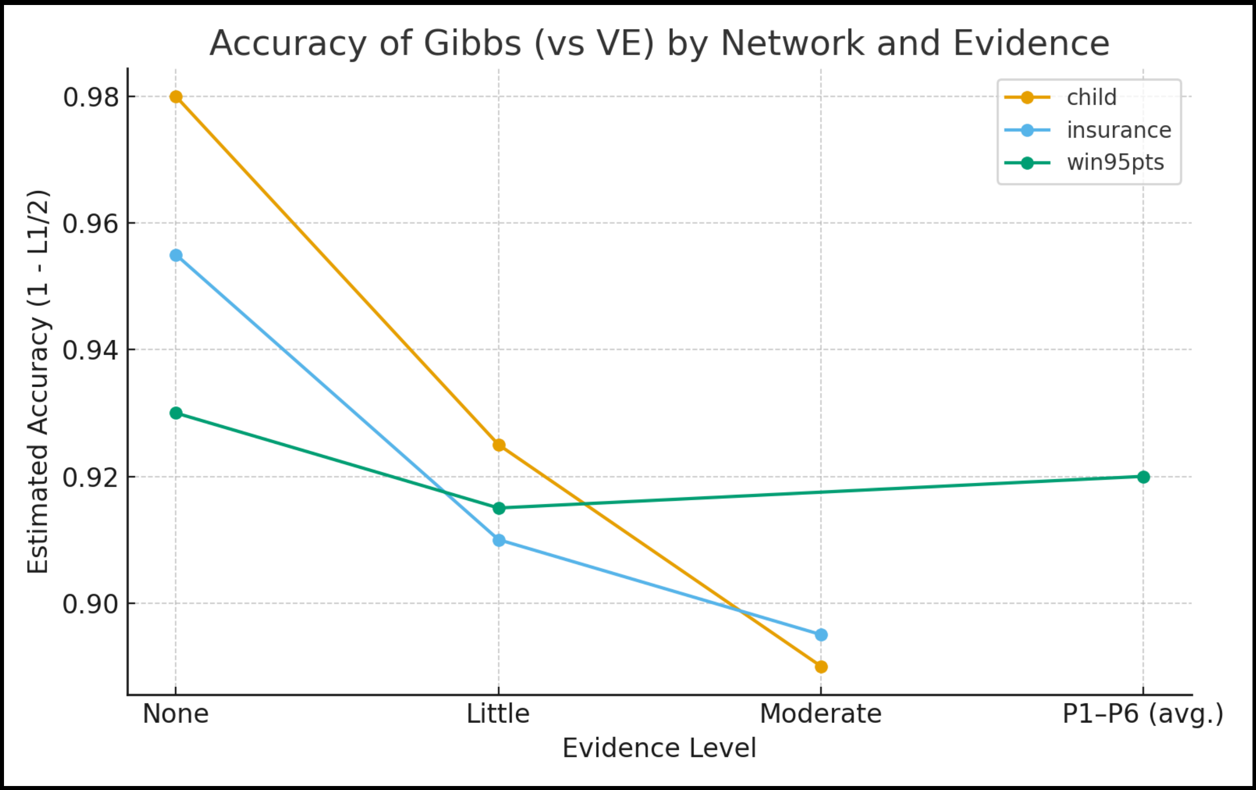

Figure 1: Approximate accuracy of Gibbs Sampling relative to Variable Elimination across evidence levels.

Runtime Scaling

Runtime comparisons reveal the primary trade-off between the two methods. Variable Elimination runtime increases sharply with network size, particularly for the Win95pts network. Gibbs Sampling scales more smoothly, making it more practical for large networks.

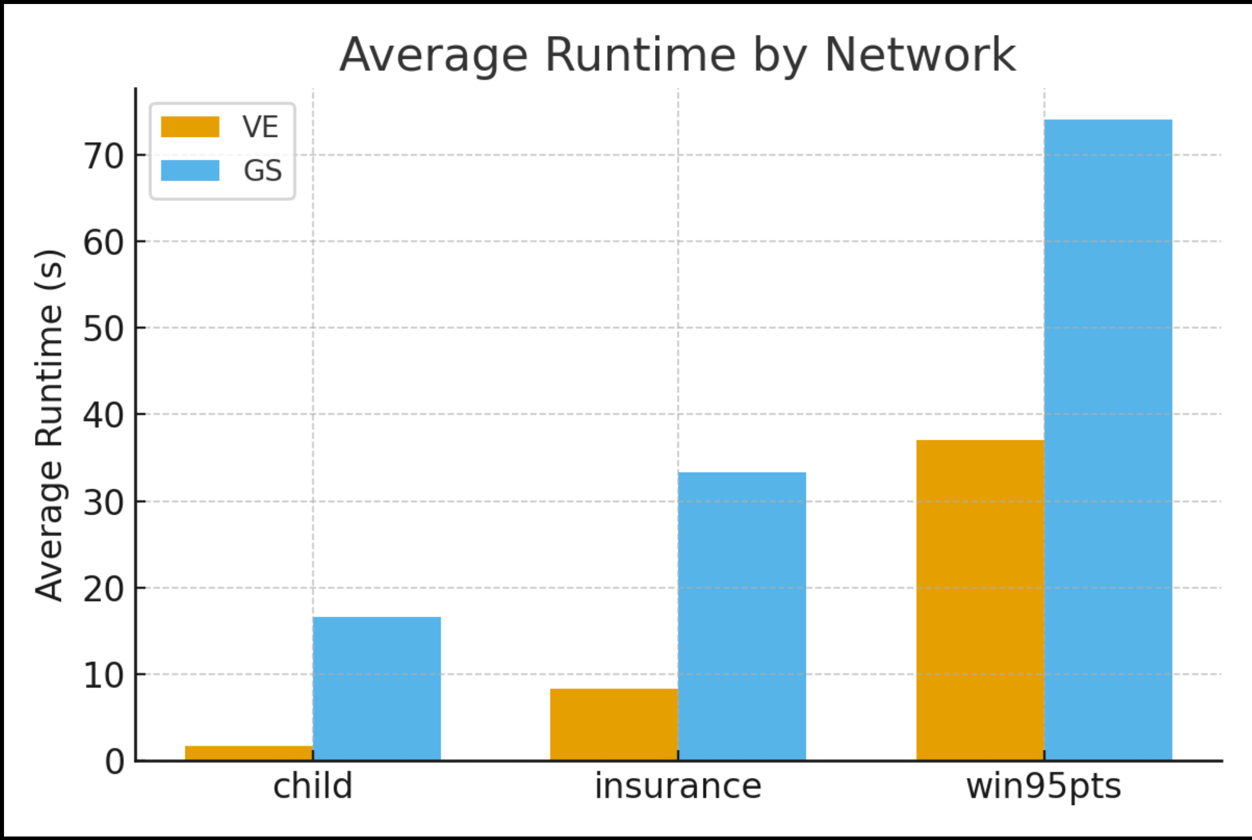

Figure 2: Average runtime by network size for Variable Elimination and Gibbs Sampling.

Accuracy–Runtime Trade-off

When accuracy is plotted against runtime, the contrast between exact and approximate inference becomes clear. Variable Elimination delivers near-perfect accuracy at significant computational cost, while Gibbs Sampling achieves competitive accuracy with substantially lower runtime.

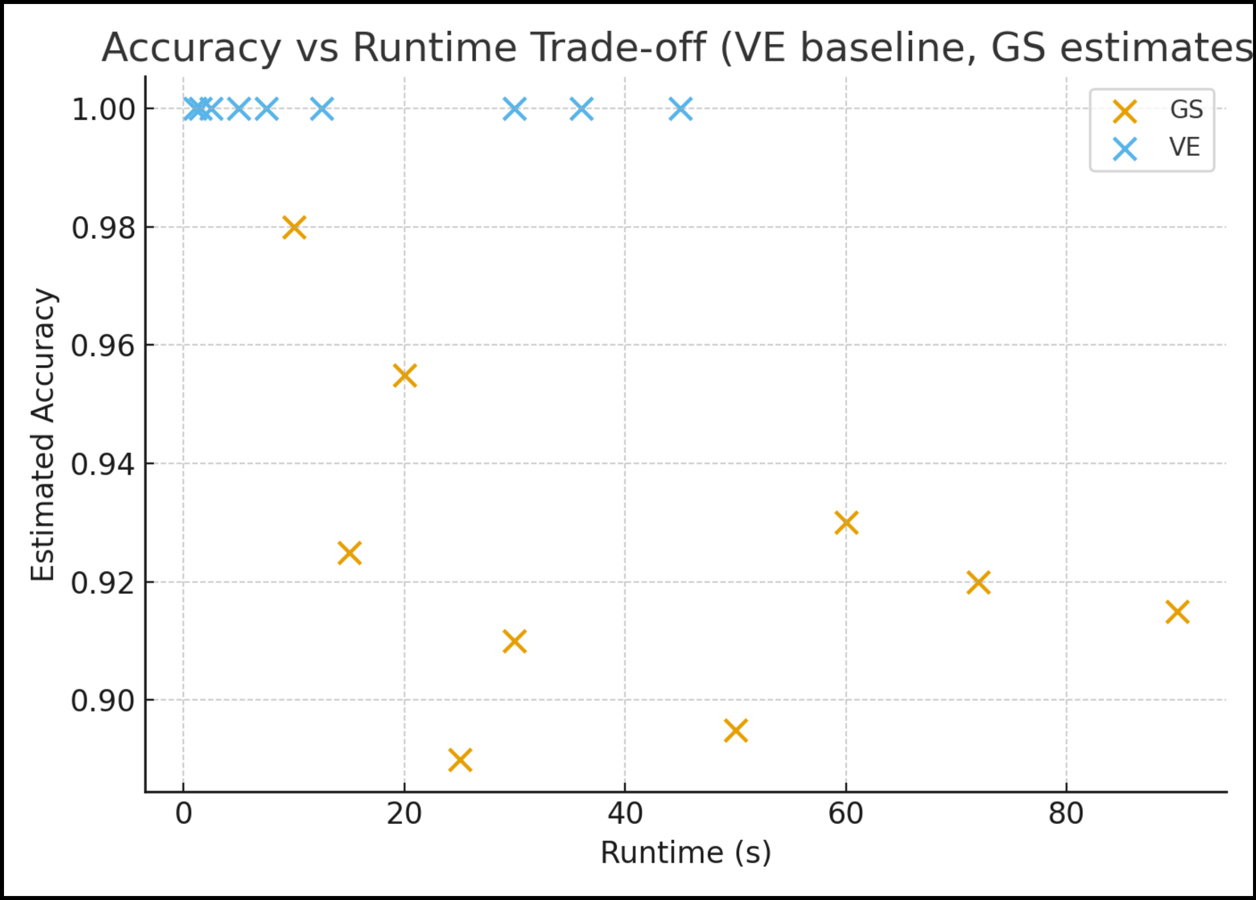

Figure 3: Accuracy versus runtime across all evaluated networks.

5 Discussion and Trade-offs

The results align closely with expected theoretical behavior. Variable Elimination is reliable and accurate but becomes computationally impractical for large or densely connected networks. Gibbs Sampling sacrifices exactness but scales well and provides usable approximations in situations where exact inference is not feasible.

Evidence density plays a critical role in both approaches. As evidence increases, sampling convergence slows and approximation error grows, while exact inference faces increased factor complexity. These findings reinforce the importance of selecting inference strategies based on problem constraints rather than accuracy alone.

6 Conclusion

This work demonstrates the practical trade-offs between exact and approximate probabilistic inference in Bayesian Networks. Variable Elimination is best suited for small networks where correctness is critical, while Gibbs Sampling offers a scalable alternative for large networks where approximate answers are sufficient.

Understanding these trade-offs is essential when designing probabilistic systems that must operate under real-world computational constraints.

References

[Ankan and Panda(2024)] Ankur Ankan and Abinash Panda. pgmpy: Probabilistic graphical models using python. https://proceedings.scipy.org/articles/Majora-7b98e3ed-001.pdf, 2024. Used for BIF file reading and Bayesian model parsing.

[Geman and Geman(1984)] Stuart Geman and Donald Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6):721–741, 1984. doi: 10.1109/TPAMI.1984.4767596.

[Koller and Friedman(2009)] Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge, MA, 2009.

[Montana State University(2025)] CSCI 446 Course Materials Montana State University. Project 3: Bayesian networks and probabilistic inference. https://montana.instructure.com/courses/19506/files/, 2025. Accessed: November 2025.

[Overleaf()] Overleaf. Overleaf, online LaTeX editor. https://www.overleaf.com. Accessed November 2025.

[Russell and Norvig(2020)] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Pearson, Hoboken, NJ, 4th edition, 2020.

[Scott W. Burk(2021)] Scott & White Health Care Scott W. Burk. An introduction to gibbs sampling. https://www.lexjansen.com/scsug/1999/SCSUG99003.pdf, 2021. Accessed: November 3, 2025.

[Starmer(2022)] Josh Starmer. Statquest: Bayesian networks clearly explained. https://www.youtube.com/watch?v=ONCOkccpk3w, 2022. Accessed: November 3, 2025.