

# Interpretability-Performance Trade-offs in Gaussian Kernels with Full-Rank vs. Diagonal Covariance on OpenML-100

Assignee Research

June 9, 2026

## Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: What is the trade-off between model interpretability and predictive performance when comparing full-rank precision matrix Gaussian kernels to diagonal covariance approximations on standardized. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: The Shrinkage-Delinkage Trade-off: An Analysis of Factorized Gaussian Approximations for Variational Inference. Research question: What is the trade-off between model interpretability and predictive performance when comparing full-rank precision matrix Gaussian kernels to diagonal covariance approximations on standardized tabular benchmarks like OpenML-100?.

## 2 Methodology

Systematic literature search across multiple databases yielded 13 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.5/10.

## 3 Results

13 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.5/10.

## 4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

## References

- <http://arxiv.org/abs/1212.0325v1>
- <http://arxiv.org/abs/2504.20900v1>
- <http://arxiv.org/abs/2302.09163v3>