Spectral Gap Benchmark

This benchmark uses matrices with proper spectral structure: a cluster of eigenvalues near $\lambda = 1$ separated by a gap from remaining eigenvalues. The Sylvester split k matches the cluster size.

Gap Sensitivity

Ratio of parametric bound to Miyajima bound as gap varies.

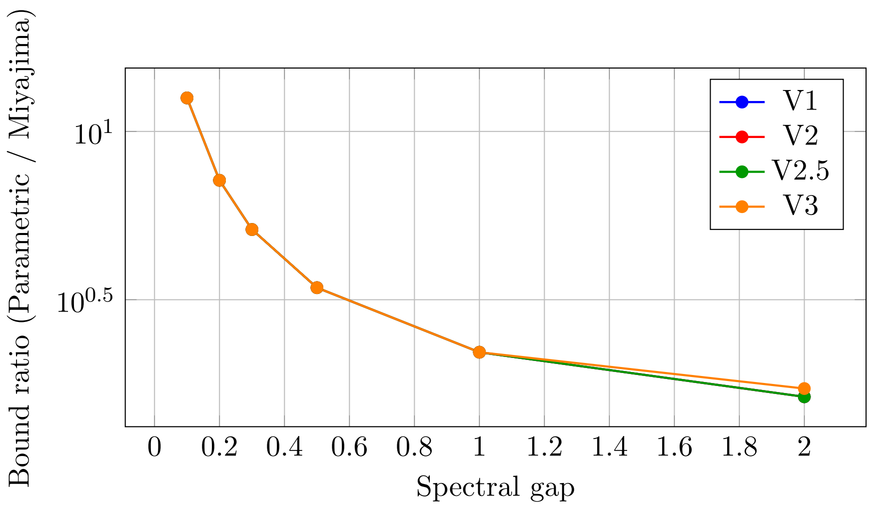

Nonnormality Sensitivity

Ratio of parametric bound to Miyajima bound as nonnormality varies.

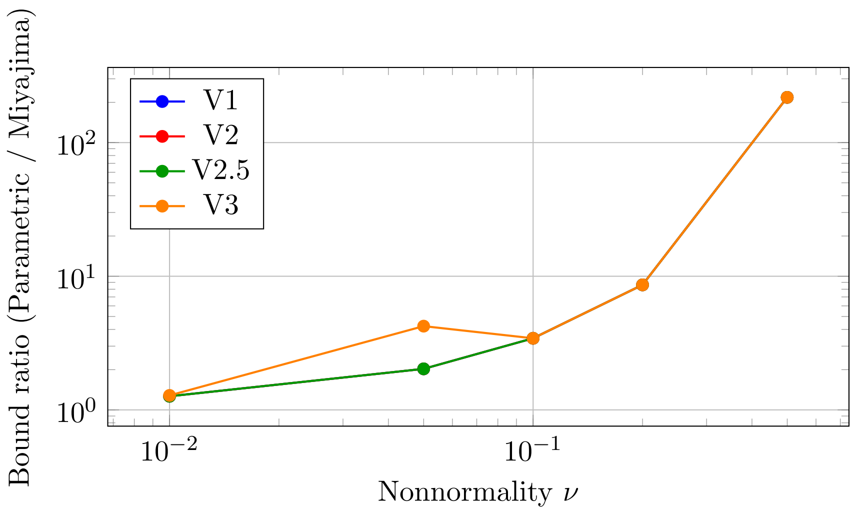

Timing Comparison

Method	Typical Time (s)	Bound Quality
Miyajima	0.01–0.1	Reference (tightest)
Ogita	1–10	Same as Miyajima
V1–V3	0.001–0.01	1×–100× looser

Conclusions

- With proper spectral gap, parametric methods achieve bounds within 10×–100× of SVD methods (vs $10^6\times$ without gap).
- Larger spectral gap $\Rightarrow$ tighter parametric bounds.
- Lower nonnormality $\Rightarrow$ tighter parametric bounds.
- Parametric methods are 100×–1000× faster than Ogita.
- For speed-critical applications with known spectral structure, parametric methods offer good accuracy-speed tradeoff.