PhD Seminar Talk 1

Randomized Estimation of T-Eigenvalues of T-SPD Tensors: A Two-Sided Bracket

Hemant Sharma[1]

[1]Indian Institute of Information Technology, Design and Manufacturing Kancheepuram, Chennai, India – 600127

Abstract: How can we estimate the largest eigenvalue of a tensor when the problem is too large to compute it directly? Under the T-product of Kilmer and Martin, the T-eigenvalues of a tensor $\mathcal{A} \in \mathbb{R}^{n\times n\times p}$ are the eigenvalues of its block circulant matrix $\mathrm{bcirc}(\mathcal{A})$ of size $d = np$. For $d = 131{,}072$, storing this matrix in dense form would need about 137 GB, so a full eigendecomposition is not an option. Deterministic bounds such as the TDep bound from our earlier work cost almost nothing, but they become loose as d grows, roughly like $\sqrt{d-1}$. Randomized methods offer a middle path, and the talk is divided into two parts.

$\star$ **Randomized power and subspace methods**: Since $\mathrm{bcirc}(\mathcal{A})$ is block diagonalized by the FFT, we can multiply it with a vector cheaply without ever forming it. Using this, we apply the power method with a random starting vector. The resulting Rayleigh quotient is always a lower bound on λ_1, and it converges exponentially at a rate governed by the spectral gap. Following Kuczyński and Woźniakowski, we give a high-probability bound that accounts for the random start. When several leading eigenvalues are needed, we adapt the randomized subspace iteration of Halko, Martinsson and Tropp to the T-product setting and prove an error bound for it. The bound makes clear how each additional power step damps the effect of the remaining spectrum. In our experiments, an oversampling of $\ell = 5$ with $q = 2$ power steps already gives less than 2% error in λ_1, and the method runs up to 67 times faster than a full eigendecomposition.

$\star\star$ **A two-sided bracket and an application to a PDE problem**: The randomized estimate never exceeds λ_1, while the TDep bound $m + s\sqrt{d-1}$ never falls below it. Putting the two together gives an interval $[\hat{\lambda}_1^{\mathrm{pow}},\ \hat{\lambda}_1^{\mathrm{TDep}}]$ that is guaranteed to contain the largest T-eigenvalue, and it did so in every trial of our validation. We also discuss a variant in which the traces are estimated with the Hutchinson estimator, which is useful when only matrix–vector products are available. As an application, we consider the discrete 3D Laplacian on a periodic slab, a simplified model of heat conduction in layered geothermal media, with sizes up to $d = 131{,}072$. When the coefficients vary, the FFT structure is lost, and the usual choice is sparse Lanczos; here the randomized power method is 4 to 29 times faster than `eigsh`. Finally, we use the estimate inside a Chebyshev iterative solver. With a 10% safety margin, the solver needs only about 3% more iterations than with the exact eigenvalue, whereas using a slight under-estimate of λ_1 makes it diverge. This example shows why a guaranteed upper bound is important in practice.

Keywords: T-product, T-eigenvalue, Randomized numerical linear algebra, Power method, Subspace iteration, Hutchinson estimator, Two-sided bracket.