Randomly pivoted Cholesky

Joel Tropp
Caltech, USA
jtropp@caltech.edu

Kernel methods are used for prediction and clustering in many data science and scientific computing applications, but applying kernel methods to a large number of data points N is expensive due to the high cost of manipulating the $N \times N$ kernel matrix. A basic approach for speeding up kernel computations is low-rank approximation, in which we replace the kernel matrix $\mathbf{A}$ with a factorized approximation that can be stored and manipulated more cheaply. When the kernel matrix $\mathbf{A}$ has rapidly decaying eigenvalues, mathematical existence proofs guarantee that $\mathbf{A}$ can be accurately approximated using a constant number of columns (without ever looking at the full matrix). Nevertheless, for a long time, designing a practical and provably justified algorithm to select the appropriate columns proved challenging.

This talk introduces Randomly Pivoted Cholesky (RPC), a natural algorithm for approximating an $N \times N$ positive-semidefinite matrix using k adaptively sampled columns. RPC can be implemented with just a few lines of code; it requires only $(k+1)N$ entry evaluations and $\mathcal{O}(k^2 N)$ additional arithmetic operations. In experiments, RPC matches or improves on the performance of alternative algorithms for low-rank psd approximation. Moreover, RPC provably achieves near-optimal approximation guarantees. The simplicity, effectiveness, and robustness of this algorithm strongly support its use for large-scale kernel computations. This work offers an accessible example of the power of using randomized algorithms for linear algebra computations.

Available at arXiv:2207.06503.

Joint work with Yifan Chen (Caltech), Ethan Epperly (Caltech) and Rob Webber (Caltech).