Accepted Manuscript

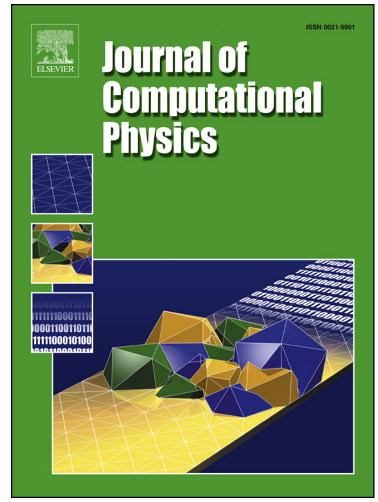

Estimating the trace of the matrix inverse by interpolating from the diagonal of an approximate inverse

Lingfei Wu, Jesse Laeuchli, Vassilis Kalantzis, Andreas Stathopoulos, Efstratios Gallopoulos

PII: S0021-9991(16)30412-0
DOI: http://dx.doi.org/10.1016/j.jcp.2016.09.001
Reference: YJCPH 6811

To appear in: *Journal of Computational Physics*

Received date: 16 June 2015
Revised date: 1 August 2016
Accepted date: 1 September 2016

Estimating the Trace of the Matrix Inverse by Interpolating from the Diagonal of an Approximate Inverse

Lingfei Wu[a,*], Jesse Laeuchli[a], Vassilis Kalantzis[b], Andreas Stathopoulos[a,*], Efstratios Gallopoulos[c]

[a]Department of Computer Science, College of William and Mary, Williamsburg, VA 23187, United States
[b]Department of Computer Science, University of Minnesota, Minneapolis, MN 55455, United States
[c]Department of Computer Engineering and Informatics, University of Patras, Patras, Greece

Abstract

A number of applications require the computation of the trace of a matrix that is implicitly available through a function. A common example of a function is the inverse of a large, sparse matrix, which is the focus of this paper. When the evaluation of the function is expensive, the task is computationally challenging because the standard approach is based on a Monte Carlo method which converges slowly. We present a different approach that exploits the pattern correlation, if present, between the diagonal of the inverse of the matrix and the diagonal of some approximate inverse that can be computed inexpensively. We leverage various sampling and fitting techniques to fit the diagonal of the approximation to the diagonal of the inverse. Depending on the quality of the approximate inverse, our method may serve as a standalone kernel for providing a fast trace estimate with a small number of samples. Furthermore, the method can be used as a variance reduction method for Monte Carlo in some cases. This is decided dynamically by our algorithm. An extensive set of experiments with various technique combinations on several matrices from some real applications demonstrate the potential of our method.

Keywords: Matrix trace, Monte Carlo method, variance reduction, preconditioner, fitting, interpolation

1. Introduction

Computing the trace of a matrix A that is given explicitly is a straightforward operation. However, for numerous applications we need to compute the trace of a matrix that is given implicitly by its action on a vector x, i.e., Ax. Specifically, many applications are interested in computing the trace of a function of a matrix $F(A)$. Examples include estimating parameters in image restoration using the generalized cross-validation approach [1], exploring the inverse covariance matrix in uncertainty quantification [2, 3], computing observables in lattice quantum chromodynamics (LQCD) [4], or counting triangles in large graphs [5]. The matrix A is large, and often sparse, so its action $F(A)x$ is typically computed through iterative methods. Because it is challenging to compute $F(A)$ explicitly, the Monte Carlo (MC) approach has become the standard method for computing the trace by averaging samples of the bilinear form $x^T F(A)x$ [6, 7]. The main purpose of this paper is to develop practical numerical techniques to address the computation of the trace of the inverse of a large, sparse matrix. But our technique can also be adapted to other functions such as the trace of the logarithm (yielding the determinant) or the trace of the matrix exponential.

For small size problems, computing A^{-1} through a dense or sparse LDU decomposition is the most efficient and accurate approach [8, 9]. This works well for discretizations of differential operators in low dimensions but becomes intractable in high dimensional discretizations. For larger size problems, domain

*Corresponding author
Email addresses: lwu@email.wm.edu (Lingfei Wu), andreas@cs.wm.edu (Andreas Stathopoulos)