Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 108C (2017) 755–764

Procedia Computer Science

International Conference on Computational Science, ICCS 2017, 12-14 June 2017, Zurich, Switzerland

TNT-NN: A Fast Active Set Method for Solving Large Non-Negative Least Squares Problems

J.M. Myre[1], E. Frahm[2], D. J. Lilja[3], and M.O. Saar[4]

[1] Department of Computer and Information Sciences, University of St. Thomas, Saint Paul, MN
myre@stthomas.edu

[2] School of Physics and Astronomy, University of Minnesota, Twin Cities
frahm@physics.umn.edu

[3] Department of Electrical and Computer Engineering, University of Minnesota, Twin Cities
lilja@umn.edu

[4] Department of Earth Sciences, ETH-Zürich
saarm@ethz.ch

Abstract

In 1974 Lawson and Hanson produced a seminal active set strategy to solve least-squares problems with non-negativity constraints that remains popular today. In this paper we present TNT-NN, a new active set method for solving non-negative least squares (NNLS) problems. TNT-NN uses a different strategy not only for the construction of the active set but also for the solution of the unconstrained least squares sub-problem. This results in dramatically improved performance over traditional active set NNLS solvers, including the Lawson and Hanson NNLS algorithm and the Fast NNLS (FNNLS) algorithm, allowing for computational investigations of new types of scientific and engineering problems.

For the small systems tested (5000×5000 or smaller), it is shown that TNT-NN is up to $95\times$ faster than FNNLS. Recent studies in rock magnetism have revealed a need for fast NNLS algorithms to address large problems (on the order of $10^5 \times 10^5$ or larger). We apply the TNT-NN algorithm to a representative rock magnetism inversion problem where it is $60\times$ faster than FNNLS. We also show that TNT-NN is capable of solving large (45000×45000) problems more than $150\times$ faster than FNNLS. These large test problems were previously considered to be unsolvable, due to the excessive execution time required by traditional methods.

Peer-review under responsibility of the scientific committee of the International Conference on Computational Science

Keywords: non-negative least squares, active set, preconditioned conjugate gradient, rock magnetism

1 Introduction

The least squares method produces the solution to a system of linear equations that minimizes error. Although it can be applied to under-determined systems of equations, it is natural to apply the least squares method to over-determined systems that cannot be solved in a

1877-0509 © 2017 The Authors. Published by Elsevier B.V.
Peer-review under responsibility of the scientific committee of the International Conference on Computational Science
10.1016/j.procs.2017.05.194

way that will satisfy all of the equations. Since the inception of the least squares method by Gauss [3], it has been applied in countless scientific, engineering, and numerical contexts. When investigating physical systems, it is not uncommon for non-negativity constraints to be applied in order for the system to be physically meaningful. Such constraints have been necessary in the computational Earth sciences when using nuclear magnetic resonance to determine pore structures [7], formulating petrologic mixing models [1, 21], and addressing seismologic inversion problems [28]. This constrained least squares problem is the non-negative least-square (NNLS) problem, that can be stated as $min_{\boldsymbol{x}}||\boldsymbol{Ax}-\boldsymbol{b}||_2$, such that $\boldsymbol{x} \geq 0$, where $\boldsymbol{A}$ is the $m \times n$ system of equations, $\boldsymbol{b}$ is the solution vector of measured data, and $\boldsymbol{x}$ is the vector of obtained parameters obtained that minimizes the 2-norm of the residual.

Recent problems in rock magnetism have required solving NNLS problems composed of tens of thousands of variables [8, 26, 25]. Without the restrictions of conventional computer systems and algorithms, these problems could expand to hundreds of thousands of variables. Direct spatial solutions to these problems using standard NNLS algorithms have required months of execution time. Frequency domain methods [15] have been developed as an alternative fast method to obtain representative solutions, but can produce non-physical artifacts (*e.g.,* Gibbs phenomenon at sharp interfaces) that violate the non-negativity constraint.

We present TNT-NN, a new (dynamite) active set strategy for solving large NNLS problems. TNT-NN improves upon prior efforts by incorporating a more aggressive strategy for identifying the active set of constraints and by using an improved solver to address the unconstrained least squares sub-problem. We show that TNT-NN dramatically outperforms the present Fast NNLS active set algorithm on a wide variety of test problems and that the maximum tractable problem size is extended to the point where previously prohibitive problems are now feasible.

2 Background and related work

Active set strategies for the NNLS problem can be broken down into two basic parts, 1) the strategy for identifying the "active" non-negativity constraints and 2) the strategy for solving the unconstrained least squares sub-problem for each choice of the active set of constraints. These parts are independent of one another and can therefore be discussed independently.

In many applications of NNLS the maximum problem size that can be investigated is limited by the algorithmic performance and the available computational resources. One of the most widely used is the 1974 Lawson and Hanson NNLS algorithm [14] (LH-NNLS). Like most non-negative active set strategies, the LH-NNLS algorithm attempts to find the non-negative solution by setting some variables to zero. These variables are the active set, because their non-negativity constraints are "active". In each iteration, the active set is modified by a single variable. The active set variables are then ignored and an unconstrained least squares sub-problem is solved. This conservative approach can yield extremely slow convergence on problems with many variables. The LH-NNLS algorithm is found in popular software packages such as Matlab, GNU R [17], and scientific tools for python [12]. LH-NNLS is commonly encountered in reference literature as the method to solve NNLS problems [2, 20, 27] and its prevalence continues to inspire new works focused on its optimization [10, 24, 16].

The Fast NNLS (FNNLS) algorithm of Bro and De Jong [6] improves upon LH-NNLS by by avoiding redundant computations and by allowing an initial "loading" of the active set. Although the FNNLS algorithm will continue to modify the active set by one variable per iteration, the total number of iterations that are required for convergence is reduced by supplying an initial solution that is close to the actual solution. Bro and De Jong report FNNLS speedups over LH-NNLS ranging from 2x-5x using small real and synthetic test suites.