

This article was downloaded by: [Lewis, A. S.]

On: 1 April 2010

Access details: Access Details: [subscription number 920446956]

Publisher Taylor & Francis

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Optimization

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title~content=t713646500>

Randomized Hessian estimation and directional search

D. Leventhal^a; A. S. Lewis^a

^a School of Operations Research and Information Engineering, Cornell University, Ithaca, NY 14853, USA

First published on: 29 March 2010

To cite this Article Leventhal, D. and Lewis, A. S. (2010) 'Randomized Hessian estimation and directional search', *Optimization*, First published on: 29 March 2010 (iFirst)

To link to this Article: DOI: 10.1080/02331930903100141

URL: <http://dx.doi.org/10.1080/02331930903100141>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

Randomized Hessian estimation and directional search

D. Leventhal and A.S. Lewis*

*School of Operations Research and Information Engineering, Cornell University,
Ithaca, NY 14853, USA*

(Received 25 September 2008; final version received 20 May 2009)

We explore how randomization can help asymptotic convergence properties of simple directional search-based optimization methods. Specifically, we develop a cheap, iterative randomized Hessian estimation scheme. We then apply this technique and analyse how it enhances a random directional search method. Then, we proceed to develop a conjugate-directional search method that incorporates estimated Hessian information without requiring the direct use of gradients.

Keywords: derivative-free optimization; directional search; quasi-Newton; random search; steepest descent

AMS Subject Classifications: 90C56; 90C53; 65K05

1. Introduction and motivation

Stochastic techniques in directional search algorithms have been well-studied in solving optimization problems, often where the underlying functions themselves are random or noisy. Some of these algorithms are based on directional search methods that obtain a random search direction which approximates a gradient in expectation: for some background on this class of algorithms, see [4, Ch. 6] or [12, Ch. 5]. For many algorithms, the broad convergence theory, combined with inherent computational simplicity, makes them particularly appealing, even for noiseless, deterministic optimization problems.

In this work, we avoid any direct use of gradient information, relying only on function evaluations. In that respect, the methods we consider have the flavour of derivative-free algorithms. Our goal, however, is not the immediate development of a practical, competitive, derivative-free optimization algorithm: our aim is instead primarily speculative. In contrast with much of the derivative-free literature, we make several impractical assumptions that hold throughout this article. We assume that the function we seek to minimize is twice differentiable and that evaluations of

*Corresponding author. Email: aslewis@orie.cornell.edu; people.orie.cornell.edu/~aslewis

that function are reliable, cheap and accurate. Further, we assume that derivative information is neither available directly nor via automatic differentiation, but it is well-approximated by finite differencing. Again, for the purposes of speculation, we assume that any line search subproblem is relatively cheap to solve when compared to the cost of approximating a gradient. This last assumption is based on the fact that, asymptotically, the computational cost of a line search should be independent of the problem dimension, being a one-dimensional optimization problem, while the number of function evaluations required to obtain a gradient grows linearly with the problem dimension, based on previous assumptions. Within this narrow framework, we consider the question as to whether, in principle, randomization can be incorporated to help simple iterative algorithms achieve good asymptotic convergence.

Keeping this narrow framework in mind, this article is organized as follows. In the remainder of this section, we consider a randomized directional search algorithm that chooses a search direction uniformly at random from the unit sphere and apply it to convex quadratic functions, comparing convergence results with a traditional gradient descent algorithm. In Section 2, we introduce a technique of randomized Hessian estimation and prove some basic properties. In Section 3, we consider algorithmic applications of our randomized Hessian estimation method. In particular, we show how Hessian estimates can be used to accelerate the uniformly random search algorithm introduced in this section and, additionally, how randomized Hessian estimation can also be used to develop a conjugate direction-like algorithm.

To illustrate the use of randomization, consider the following basic algorithm: at each iteration, choose a search direction uniformly at random on the unit sphere and perform an exact line search. Note that this algorithm has been widely studied, with analysis appearing in [5] and [11], among others. Further, it was shown to be linearly convergent for twice differentiable functions under conditions given in [10].

Consider applying this algorithm to the problem of minimizing a convex quadratic function $f(x) = \frac{1}{2}x^T Ax + b^T x$ where A is a symmetric, positive-definite, $n \times n$ matrix. First, defining the energy norm by $\|x\|_A = \sqrt{x^T Ax}$, observe that if the unique minimizer of f is $x^* = -A^{-1}b$, then the identity

$$f(x) - f(x^*) = \frac{1}{2}(x - x^*)^T A(x - x^*) = \frac{1}{2}\|x - x^*\|_A^2 \quad (1.1)$$

holds. Observe that if the current iterate is x , then the new iterate is given by

$$x_+ = x - \frac{d^T(Ax + b)}{d^T Ad} d \quad (1.2)$$

and the new function value is

$$f(x_+) = f(x) - \frac{(d^T(Ax + b))^2}{2d^T Ad}.$$

The difference between the current function value and the optimal value is reduced by the ratio

$$\begin{aligned} \frac{f(x_+) - f(x^*)}{f(x) - f(x^*)} &= 1 - \frac{(d^T(Ax + b))^2}{2(d^T Ad)(f(x) - f(x^*))} \\ &= 1 - \frac{(d^T(Ax + b))^2}{(d^T Ad)((x - x^*)^T A(x - x^*))} \\ &= 1 - \frac{(d^T A(x - x^*))^2}{(d^T Ad)((A(x - x^*))^T A^{-1}(A(x - x^*)))} \\ &\leq 1 - \frac{1}{\kappa(A)} \left(d^T \frac{A(x - x^*)}{\|A(x - x^*)\|} \right)^2, \end{aligned}$$

where $\kappa(A) = \|A\| \|A^{-1}\|$ denotes the condition number of A . Observe that the distribution of d is invariant under orthogonal transformations. Therefore, let U be any orthogonal transformation satisfying $U \left(\frac{A(x - x^*)}{\|A(x - x^*)\|} \right) = e_1$, the first standard basis vector. From this, we have

$$\begin{aligned} \mathbf{E} \left[\left(d^T \frac{A(x - x^*)}{\|A(x - x^*)\|} \right)^2 \mid x \right] &= \mathbf{E} \left[\left((U^T d)^T \frac{A(x - x^*)}{\|A(x - x^*)\|} \right)^2 \mid x \right] \\ &= \mathbf{E}[d_1^2] \\ &= \frac{1}{n} \mathbf{E} \left[\left(\sum_i d_i^2 \right) \right] \\ &= \frac{1}{n}, \end{aligned}$$

where the first equality follows from the invariance of the distribution of d and the third equality follows from the fact that each component of d is identically distributed. We deduce

$$\mathbf{E}[f(x_+) - f(x^*) \mid x] \leq \left(1 - \frac{1}{n\kappa(A)} \right) (f(x) - f(x^*)) \tag{1.3}$$

with equality when A is a multiple of the identity matrix, in which case $\kappa(A) = 1$.

Compare this with the steepest descent algorithm. A known result about the steepest descent algorithm in [1] says that given an initial iterate x and letting \hat{x} be the new iterate constructed from an exact line search in the negative gradient direction,

$$f(\hat{x}) - f(x^*) \leq \left(\frac{\kappa(A) - 1}{\kappa(A) + 1} \right)^2 (f(x) - f(x^*)) = \left(1 - O\left(\frac{1}{\kappa(A)} \right) \right) (f(x) - f(x^*)).$$

Further, for most initial iterates x , this inequality is asymptotically tight if this procedure is iteratively repeated. Consider the following asymptotic argument, applying the assumptions made earlier in this section. Supposing derivative information is only available through – and well-approximated by – finite differencing then we can perform an exact (or almost-exact) line search in some constant number, $O(1)$, of function evaluations. It follows that each iteration of random search takes $O(1)$ function evaluations. However, since derivative

information is only available via finite differencing, computing a gradient takes $O(n)$ function evaluations. Letting \bar{x} be the iterate after performing $O(n)$ iterations of random search, we obtain that

$$E \left[\frac{f(\bar{x}) - f(x^*)}{f(x) - f(x^*)} | x \right] \leq \left(1 - \frac{1}{n\kappa(A)} \right)^{O(n)} = 1 - O\left(\frac{1}{\kappa(A)} \right).$$

Essentially, the expected improvement of random search is on the same order of magnitude as steepest descent when measured on a cost per function evaluation basis. This simple example suggests that randomization techniques may be an interesting ingredient in the design and analysis of iterative optimization algorithms.

2. Randomized Hessian estimation

In this section, we will consider arbitrary twice-differentiable functions $f: \mathbf{R}^n \rightarrow \mathbf{R}$. As in the previous section, assume that these functions can be evaluated exactly, but derivative information is only available through finite differencing. In particular, for any vector $v \in \mathbf{R}^n$, suppose we can use finite differencing to well-approximate the second derivative of f at x in the direction v via the formula

$$v^T \nabla^2 f(x) v \approx \frac{f(x + \epsilon v) - 2f(x) + f(x - \epsilon v)}{\epsilon^2} \quad (2.1)$$

for some sufficiently small $\epsilon > 0$. In particular, note that by choosing $\frac{1}{2}n(n+1)$ suitable directions v , we could effectively approximate the entire Hessian $\nabla^2 f(x)$.

In Section 1, we considered a framework in which computational costs of an algorithm are measured by the number of function evaluations required and we will continue with that throughout this article. In particular, it was shown that under this framework, the steepest descent algorithm, asymptotically, achieves improvement on the same order of magnitude as a uniformly random search algorithm when applied to convex quadratics. Ideally, we would like to extend these methods of analysis to algorithms that incorporate additional information about a function's behaviour. For example, instead of calculating a complete Hessian matrix at each iteration, Newton-like methods rely on approximations to the Hessian matrix which are iteratively updated, often from successively generated gradient information. To consider a similar approach in the context of random search, suppose we begin with an approximation to the Hessian matrix, denoted B , and some unit vector $v \in \mathbf{R}^n$. Consider the new matrix B_+ obtained by making a rank-one update so that the new matrix B_+ matches the true Hessian in the direction v , i.e.

$$B_+ = B + (v^T(\nabla^2 f(x) - B)v)vv^T. \quad (2.2)$$

This rank-one update results in the new matrix B_+ having the property that $v^T B_+ v = v^T \nabla^2 f(x) v$. Note that if this update is performed using the approximate second derivative via Equation (2.1), then this costs only three function evaluations. The following result on the space of symmetric $n \times n$ matrices, \mathbf{S}^n , equipped with the usual trace inner product and the induced Frobenius norm will be our primary tool.

THEOREM 2.1 *Given any matrices $H, B \in \mathbf{S}^n$, if the random vector $d \in \mathbf{R}^n$ is uniformly distributed on the unit sphere, then the matrix*

$$B_+ = B + (d^T(H - B)d)dd^T$$

satisfies

$$\|B_+ - H\| \leq \|B - H\|$$

and

$$\mathbf{E}[\|B_+ - H\|^2] \leq \left(1 - \frac{2}{n(n+2)}\right) \|B - H\|^2.$$

Proof Since we can rewrite the update in the form

$$(B_+ - H) = (B - H) - (d^T(B - H)d)dd^T,$$

we lose no generality in assuming $H=0$. Again, we lose no generality in assuming $\|B\| = 1$, and proving

$$\|B_+\| \leq 1 \quad \text{and} \quad \mathbf{E}[\|B_+\|^2] \leq 1 - \frac{2}{n(n+2)}.$$

From the equation

$$B_+ = B - (d^T B d)dd^T,$$

we immediately deduce

$$\|B_+\|^2 = \|B\|^2 - (d^T B d)^2 = 1 - (d^T B d)^2 \leq 1.$$

To complete the proof, we need to bound the quantity $\mathbf{E}[(d^T B d)^2]$. We can diagonalize the matrix $B = U^T(\text{Diag } \lambda)U$ where U is orthogonal and the vector of eigenvalues $\lambda \in \mathbf{R}^n$ satisfies $\|\lambda\| = 1$ by assumption. Using the fact that the distribution of d is invariant under orthogonal transformations, we obtain

$$\begin{aligned} \mathbf{E}[(d^T B d)^2] &= \mathbf{E}[(d^T U^T(\text{Diag } \lambda)U d)^2] = \mathbf{E}[(d^T(\text{Diag } \lambda)d)^2] \\ &= \mathbf{E}\left[\left(\sum_{i=1}^n \lambda_i d_i^2\right)^2\right] = \mathbf{E}\left[\sum_i \lambda_i^2 d_i^4 + \sum_{i \neq j} \lambda_i \lambda_j d_i^2 d_j^2\right] \\ &= \mathbf{E}[d_1^4] + \left(\sum_{i \neq j} \lambda_i \lambda_j\right) \mathbf{E}[d_1^2 d_2^2] \end{aligned}$$

by symmetry. Since we know that

$$0 \leq \left(\sum_i \lambda_i\right)^2 = \sum_i \lambda_i^2 + \sum_{i \neq j} \lambda_i \lambda_j = 1 + \sum_{i \neq j} \lambda_i \lambda_j,$$

it follows that

$$\mathbf{E}[(d^T B d)^2] \geq \mathbf{E}[d_1^4] - \mathbf{E}[d_1^2 d_2^2].$$

Standard results on integrals over the unit sphere in \mathbf{R}^n gives the formula

$$\int_{\|x\|=1} x_1^\nu d\sigma = 2\pi^{\frac{n-1}{2}} \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma\left(\frac{\nu+n}{2}\right)},$$

where $d\sigma$ denotes an $(n-1)$ -dimensional surface element and Γ denotes the Gamma function. We deduce

$$\mathbf{E}(d_1^4) = \frac{\int_{\|x\|=1} x_1^4 d\sigma}{\int_{\|x\|=1} d\sigma} = \frac{\Gamma\left(\frac{5}{2}\right)}{\Gamma\left(\frac{n}{2} + 2\right)} \cdot \frac{\Gamma\left(\frac{n}{2}\right)}{\Gamma\left(\frac{1}{2}\right)} = \frac{\frac{3}{2} \cdot \frac{1}{2}}{\left(\frac{n}{2} + 1\right) \cdot \frac{n}{2}} = \frac{3}{n(n+2)}.$$

Furthermore,

$$1 = \left(\sum_i d_i^2 \right)^2 = \sum_i d_i^4 + \sum_{i \neq j} d_i^2 d_j^2,$$

so using symmetry again shows

$$1 = n\mathbf{E}[d_1^4] + n(n-1)\mathbf{E}[d_1^2 d_2^2].$$

From this we deduce

$$\mathbf{E}[d_1^2 d_2^2] = \frac{1 - n\mathbf{E}[d_1^4]}{n(n-1)} = \frac{1}{n(n+2)}.$$

Therefore, this shows that

$$\mathbf{E}[(d^T B d)^2] \geq \frac{3}{n(n+2)} - \frac{1}{n(n+2)} = \frac{2}{n(n+2)},$$

so

$$\mathbf{E}[\|B_+\|^2] \leq 1 - \frac{2}{n(n+2)}$$

as required. ■

To continue, note that iterating this procedure generates a random sequence of Hessian approximations which, as shown further, converges almost surely to the true Hessian.

COROLLARY 2.2 *Given any matrices $H, B_0 \in \mathbf{S}^n$, consider the sequence of matrices $B_k \in \mathbf{S}^n$ for $k=0, 1, 2, \dots$, defined iteratively by*

$$B_{k+1} = B_k + \left((d^k)^T (H - B_k) d^k \right) d^k (d^k)^T,$$

where the random vectors $d^0, d^1, d^2, \dots \in \mathbf{R}^n$ are independent and uniformly distributed on the unit sphere. Then the errors $\|B_k - H\|$ decrease monotonically, and $B_k \rightarrow H$ almost surely.

Proof Again we lose no generality in assuming $H=0$. By the previous result, it follows that the random variables $\|B_k\|^2$ form a supermartingale which is bounded

above by $\|B_0\|^2$. By classical supermartingale convergence results (see e.g. [2]), it follows that $\|B_k\|^2 \rightarrow Y$ almost surely for some random variable $Y \geq 0$. It remains to be shown that $Y=0$ almost everywhere.

Clearly, the random variables B_k are uniformly bounded. Furthermore, if we define a measure on the space of symmetric $n \times n$ matrices, γ_k , by

$$\gamma_k(U) = \text{pr}\{B_k \in U\}$$

for any measurable set U , then the previous result implies

$$\begin{aligned} \mathbf{E}[\|B_{k+1}\|^2] &= \int \mathbf{E}[\|B_{k+1}\|^2 | B_k = X] d\gamma_k(X) \\ &\leq \int \left(1 - \frac{2}{n(n+2)}\right) \|X\|^2 d\gamma_k(X) \\ &= \left(1 - \frac{2}{n(n+2)}\right) \mathbf{E}[\|B_k\|^2]. \end{aligned}$$

Now by monotone convergence we have

$$\mathbf{E}[Y] = \mathbf{E}[\lim_k \|B_k\|^2] = \lim_k \mathbf{E}[\|B_k\|^2] = 0,$$

so $Y=0$ almost everywhere as required. ■

In a more realistic framework for optimization, we wish to approximate a limiting Hessian. In the context of randomized algorithms, such as the ‘random search’ algorithm described in Section 1, the iterates generated by the algorithm are now random. By using Hessian information at each iterate to update our Hessian estimate, we now have to consider that the corresponding sequence of Hessians used is now itself random, though ideally approaching a limiting Hessian, in addition to considering the random sequence of Hessian estimates generated by the estimation procedure of Theorem 2.1. To account for this in the following theorem, recall that $\sqrt{\mathbf{E}[\|X\|^2]}$ is a norm on the space of random matrices. Applying properties of norms to this function, as the next result shows, we obtain convergence of the random Hessian estimates to the limiting Hessian.

THEOREM 2.3 *Consider a sequence of random matrices $H_k \in \mathbf{S}^n$ for $k=1, 2, 3, \dots$, with each $\mathbf{E}[\|H_k\|^2]$ finite, and a fixed matrix $\bar{H} \in \mathbf{S}^n$ such that $\mathbf{E}[\|H_k - \bar{H}\|^2] \rightarrow 0$. Consider a sequence of random matrices $B_k \in \mathbf{S}^n$ for $k=0, 1, 2, \dots$, with $\mathbf{E}[\|B_0\|^2]$ finite, related by the iterative formula*

$$B_{k+1} = B_k + \left((d^k)^T (H_k - B_k) d^k \right) d^k (d^k)^T,$$

where the random vectors $d^0, d^1, d^2, \dots \in \mathbf{R}^n$ are independent and uniformly distributed on the unit sphere. Then $\mathbf{E}[\|B_k - \bar{H}\|^2] \rightarrow 0$.

Proof By Corollary 2.2, we know for each $k=0, 1, 2, \dots$ the inequality

$$\|B_{k+1} - H_k\|^2 \leq \|B_k - H_k\|^2$$

holds. Hence by induction it follows that $\mathbf{E}[\|B_k\|^2]$ is finite for all $k \geq 0$.

Define a number

$$r = \sqrt{1 - \frac{2}{n(n+2)}} \in (0, 1).$$

By Theorem 2.1, we have

$$\mathbf{E}[\|B_{k+1} - H_k\|^2 | B_k, H_k] \leq r^2 \|B_k - H_k\|^2.$$

Once again, define a measure γ_k by

$$\gamma_k(S) = \text{pr}\{(B_k, H_k) \in S\}$$

for any measurable set S . Then we have

$$\begin{aligned} \mathbf{E}[\|B_{k+1} - H_k\|^2] &= \int \mathbf{E}[\|B_{k+1} - H_k\|^2 | (B_k, H_k) = (B, H)] d\gamma_k(B, H) \\ &\leq \int r^2 \|B - H\|^2 d\gamma_k(B, H) \\ &= r^2 \mathbf{E}[\|B_k - H_k\|^2]. \end{aligned}$$

Applying the triangle inequality property of norms gives

$$\sqrt{\mathbf{E}[\|B_{k+1} - \bar{H}\|^2]} \leq r \sqrt{\mathbf{E}[\|B_k - \bar{H}\|^2]} + (1+r) \sqrt{\mathbf{E}[\|H_k - \bar{H}\|^2]}.$$

Now fix any number $\epsilon > 0$. By assumption, there exists an integer \bar{k} such that for all integers $k \geq \bar{k}$ we have

$$\mathbf{E}[\|H_k - \bar{H}\|^2] \leq \left(\frac{\epsilon(1-r)}{2(1+r)} \right)^2.$$

Hence, for all $k \geq \bar{k}$, we deduce

$$\sqrt{\mathbf{E}[\|B_{k+1} - \bar{H}\|^2]} \leq r \sqrt{\mathbf{E}[\|B_k - \bar{H}\|^2]} + \frac{\epsilon(1-r)}{2}.$$

For such k , if $\mathbf{E}[\|B_k - \bar{H}\|^2] \leq \epsilon^2$, then

$$\sqrt{\mathbf{E}[\|B_{k+1} - \bar{H}\|^2]} \leq \frac{\epsilon(1+r)}{2} < \epsilon,$$

whereas if $\mathbf{E}[\|B_k - \bar{H}\|^2] > \epsilon^2$, then

$$\begin{aligned} \sqrt{\mathbf{E}[\|B_{k+1} - \bar{H}\|^2]} &< r \sqrt{\mathbf{E}[\|B_k - \bar{H}\|^2]} + \frac{1-r}{2} \sqrt{\mathbf{E}[\|B_k - \bar{H}\|^2]} \\ &= \frac{1+r}{2} \sqrt{\mathbf{E}[\|B_k - \bar{H}\|^2]}. \end{aligned}$$

Consequently, $\mathbf{E}[\|B_k - \bar{H}\|^2] \leq \epsilon^2$ for all large k . Since $\epsilon > 0$ was arbitrary, the result follows. ■

3. Applications to algorithms

3.1. Random search, revisited

Return to the convex quadratic function $f(x) = \frac{1}{2}x^T Ax + b^T x$ considered in Section 1, where A is a positive definite, $n \times n$ matrix and x^* is the unique minimizer. Recall that if we consider the iterative algorithm given by Equation (1.2), letting d be a unit vector uniformly distributed on the unit sphere and letting

$$x_+ = x - \frac{d^T(Ax + b)}{d^T Ad} d,$$

then it was shown in inequality (1.3) that

$$\mathbf{E}[f(x_+) - f(x^*)|x] \leq \left(1 - \frac{1}{n\kappa(A)}\right)(f(x) - f(x^*)).$$

Now, suppose that H is a positive-definite estimate of the matrix A and consider the Cholesky factor matrix C such that $CC^T = H^{-1}$. Suppose that instead of performing an exact line search in the uniformly distributed direction d , we instead perform the line search in the direction Cd . From this we obtain

$$\begin{aligned} \frac{f(x_+) - f(x^*)}{f(x) - f(x^*)} &= 1 - \frac{(d^T C^T(Ax + b))^2}{2(d^T C^T ACd)(f(x) - f(x^*))} \\ &= 1 - \frac{(d^T C^T(Ax + b))^2}{(d^T C^T ACd)((x - x^*)^T A(x - x^*))} \\ &= 1 - \frac{(d^T (C^T A(x - x^*)))^2}{(d^T (C^T AC)d)((C^T A(x - x^*))^T (C^T AC)^{-1} (C^T A(x - x^*)))} \\ &\leq 1 - \frac{1}{\kappa(C^T AC)} \left(d^T \frac{C^T A(x - x^*)}{\|C^T A(x - x^*)\|} \right)^2, \end{aligned}$$

allowing us to conclude that

$$\mathbf{E}[f(x_+) - f(x^*)|x] \leq \left(1 - \frac{1}{n\kappa(C^T AC)}\right)(f(x) - f(x^*)).$$

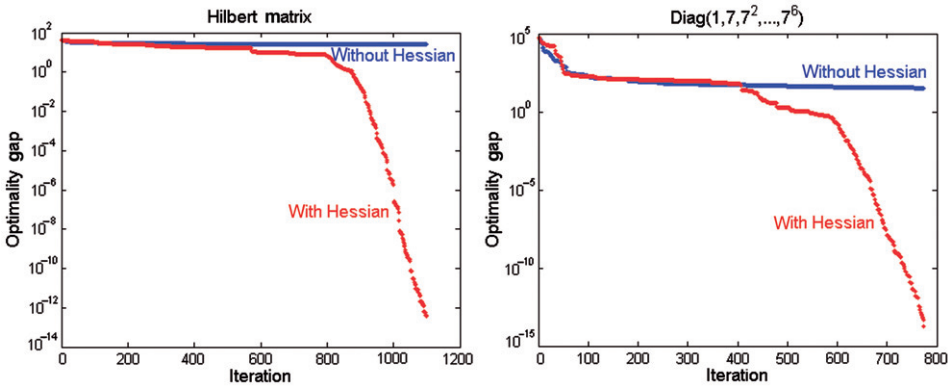
This provides the same convergence rate as performing the random search algorithm given by Equation (1.2) on the function $g(x) = \frac{1}{2}x^T (C^T AC)x + b^T x$. Consider an implementation of this algorithm using the Hessian approximation technique described in Section 2. Given a current iterate x_{k-1} and Hessian approximation B_{k-1} , we can proceed as follows. First, form the new Hessian approximation B_k given by Equation (2.2), choosing the update vector uniformly at random from the unit sphere. Observe that by Corollary 2.2, if A is positive definite, then B_k will be positive definite as well almost surely for all sufficiently large k , in which case, obtain the Cholesky factorization $B_k^{-1} = C_k C_k^T$. Otherwise, one suggested heuristic, implemented below, is to obtain the projection of B_k onto the positive semi-definite cone, denoted B_k^+ , and perform the Cholesky factorization $C_k C_k^T = (B_k^+ + \epsilon I)^{-1}$ for some $\epsilon > 0$. Finally, we can find the next iterate x_k by a line search in the direction $C_k d_k$ where d_k is uniformly distributed on the unit sphere.

Since $B_k \rightarrow A$ almost surely by Corollary 2.2, it follows $C_k \rightarrow A^{-\frac{1}{2}}$ almost surely as well. Therefore, it follows that

$$\frac{\mathbf{E}[f(x_{k+1}) - f(x^*) | x_k]}{f(x_k) - f(x^*)} \leq 1 - \frac{1}{n\kappa(C_k^T A C_k)} \rightarrow 1 - \frac{1}{n}.$$

Thus, the uniformly random search algorithm incorporating the Hessian update provides linear convergence with asymptotic rate $1 - \frac{1}{n}$, independent of the conditioning of the original matrix.

Below are two examples of the algorithm's behaviour with a convex quadratic function $f(x) = 0.5x^T A x + b^T x$, where $b = [1, 1, \dots, 1]^T$. The first example uses a Hilbert Matrix of size 7 (with condition number on the order of 10^8) while the second uses the matrix $A = \text{Diag}(1, 7, 7^2, \dots, 7^6)$. In each case, we compare uniformly the random search with the Cholesky-weighted random search described earlier, using the projection heuristic when the Hessian estimate is not positive definite. Additionally, each search vector and Hessian update vector, when applicable, was chosen independently in each example and, when applicable, an exact Hessian update was implemented.



3.2. A conjugate directions algorithm

Coordinate descent algorithms have a long and varied history in differentiable minimization. In the worst case, examples of continuously differentiable functions exist in [9] where a coordinate descent algorithm will fail to converge to a first-order stationary point. On the other hand, for twice-differentiable, strictly convex functions, variants of coordinate descent methods were shown to be linearly convergent in [6]. In either case the simplicity of such algorithms, along with the lack of a need for gradient information, often makes them appealing.

Let us briefly return to the example of a convex quadratic function $f(x) = \frac{1}{2}x^T A x + b^T x$. Consider algorithms, similar to coordinate descent algorithms, that choose search directions by cycling through some fixed set $W = \{w_1, \dots, w_n\}$, performing an exact line search at each iteration. If the search directions in W happen to be A -conjugate, satisfying $w_i^T A w_j = 0$ for all $i \neq j$, then we actually reach the optimal solution in n iterations. Alternatively, if our set of search directions fails to account for the function's second-order behaviour, convergence can be

significantly slower. Explicitly generating a set of directions that are conjugate with respect to the Hessian requires knowledge of the function's Hessian information. Methods were proposed in [8], and expanded upon in [3,7,13] among others, that begin as coordinate descent algorithms and iteratively adjust the search directions, gradually making them conjugate with respect to the Hessian matrix. Further, these adjustments are based on the results of previous line searches without actually requiring full knowledge of the Hessian or any gradients.

We propose an alternative approach for arbitrary twice-differentiable functions. If an estimate of the Hessian was readily available, we could take advantage of it by generating search directions iteratively that are conjugate with respect to the estimate. This suggests that we can design an algorithm using the Hessian estimation technique in Section 2 to dynamically generate new search directions that have the desired conjugacy properties. We can formalize this in the following algorithm.

Algorithm 3.1 Let f be a twice-differentiable function, x_0 an initial starting point, B_0 and initial Hessian estimate and $\{v_{-n}, v_{-(n-1)}, \dots, v_{-1}\}$ an initial set of search directions. For $k=0, 1, 2, \dots$

- (1) Compute the vector v_k that is B_k -conjugate to $v_{k-1}, \dots, v_{k-n+1}$.
- (2) Compute x_{k+1} as a result of a (two-way) line search in the direction v_k .
- (3) Compute B_{k+1} according to Equation (2.2), letting d_k be uniformly distributed on the unit sphere and computing

$$B_{k+1} = B_k + (d_k(\nabla^2 f(x_{k+1}) - B_k)d_k)d_k d_k^T.$$

One simple initialization scheme takes $B_0 = I$ and $\{v_{-n}, \dots, v_{-1}\} = \{e_1, \dots, e_n\}$, the standard basis vectors.

Since B_k is our Hessian approximation at the current iterate x_k , a reasonable initial step size is given by $x_{k+1} = x_k - t_k v_k$, where $t_k = \frac{v_k^T \nabla f(x_k)}{v_k^T B_k v_k}$, corresponding to an exact line search in the direction v_k of the quadratic model. The advantage to this approach is that each iteration requires only directional derivatives and, being highly iterative, this interpolates nicely with the Hessian update derived in Section 2. Specifically, when using the fixed step size mentioned earlier, each iteration takes exactly five function evaluations: $f(x_k)$, $f(x_k \pm \epsilon v_k)$, and $f(x_k \pm \epsilon d_k)$ where v_k and d_k are the search direction and the random unit vector, respectively.

The essence of this algorithm lies in using our randomized Hessian estimation technique to update a quadratic model and then performing a line search. Since we are relying solely on function evaluations, this algorithm has the 'flavour' of derivative-free optimization. However, it should be noted that a different perspective can be taken with regards to this algorithm, permitting a comparison with Newton-like methods.

Typical Newton-like methods maintain, along with the current iterate x_k , a (positive-definite) Hessian estimate B_k and proceed by performing some type of line search in the direction $-B_k^{-1} \nabla f(x_k)$. For simplicity, consider a step size of 1, i.e. $x_{k+1} = x_k - B_k^{-1} \nabla f(x_k)$. Recall that computing $B_k^{-1} \nabla f(x_k)$, equivalent to solving the system $B_k y = \nabla f(x_k)$ for y , can be done indirectly by searching in n different B_k -conjugate directions.

Specifically, suppose we have a set of directions $\{v_1, \dots, v_n\}$ that are B_k -conjugate, satisfying $v_i^T B_k v_j = 0$ for all $i \neq j$ and take $x^0 = x_k$, our current iterate. For $i = 1, \dots, n$, let $x^i = x^{i-1} - \frac{v_i^T \nabla f(x_k)}{v_i^T B_k v_i} v_i$. Then it follows that

$$x^n = x^0 - \sum_{i=1}^n \frac{v_i^T \nabla f(x_k)}{v_i^T B_k v_i} v_i = x_k - B_k^{-1} \nabla f(x_k),$$

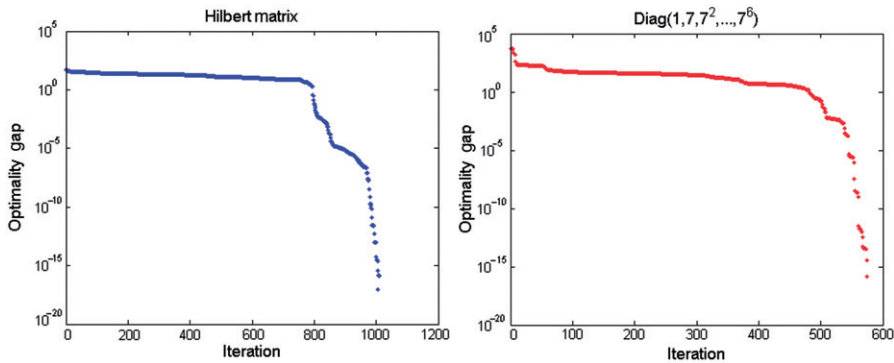
the Newton-like step. Given this interpretation of Newton-like methods, consider the version of Algorithm 3.1 where, at each iteration, the step size is fixed beforehand at $t_k = \frac{v_k^T \nabla f(x_k)}{v_k^T B_k v_k}$. Then one can interpret Algorithm 3.1 as an iterated version of a Newton-like method. Specifically, while the Newton-like method *indirectly* involves computing the quantities $v_i^T \nabla f(x_k)$ and $v_i^T B_k v_i$ with the iterate x_k and Hessian estimate B_k fixed, Algorithm 3.1 allows for a dynamically changing gradient and Hessian approximation at each conjugate direction step.

Given this connection between Algorithm 3.1 and traditional Newton-like methods, it seems natural to expect superlinear convergence under similar assumptions. This is demonstrated in the next theorem.

THEOREM 3.2 *Consider the strictly convex quadratic function $f(x) = \frac{1}{2}x^T A x + b^T x$ where A is a positive definite matrix. Then for any initial point x_0 , initial Hessian estimate B_0 and initial search directions, Algorithm 3.1 is n -step superlinearly convergent almost surely when implemented with an exact line search.*

The proof of this result can be found in the Appendix.

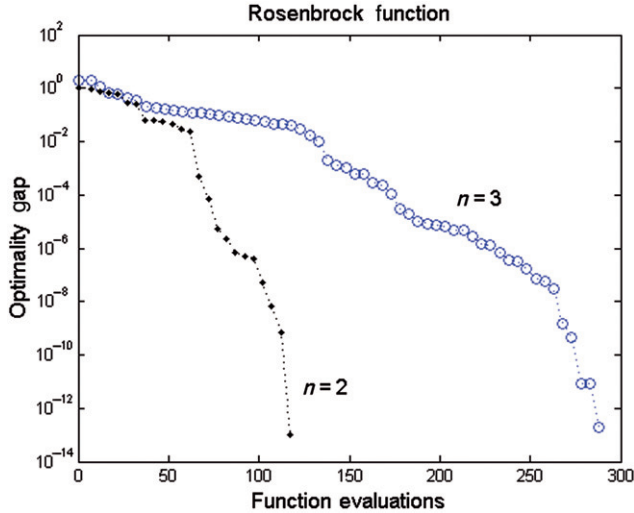
In the following example, we again consider two convex quadratic functions $\frac{1}{2}x^T A x + b^T x$ where, again, A is a Hilbert matrix of dimension 7 and $A = \text{Diag}(1, 7, 7^2, \dots, 7^6)$, respectively with $b = [1, 1, \dots, 1]^T$. The above algorithm was implemented with an exact line search and exact directional second derivatives.



Additionally, we considered two examples of this algorithm on a variant of the Rosenbrock function, given by

$$f(x) = \sum_{i=1}^{n-1} \left[(1 - x_i)^2 + 10(x_{i+1} - x_i^2)^2 \right],$$

for $n=2,3$. It was implemented with $\epsilon=10^{-4}$, initial iterate $[0,0,\dots,0]$ and a backtracking line search with initial step size equal to that suggested by the exact quadratic model, estimated via finite differencing, thereby making each iteration require five function evaluations plus any extra cost incurred by the line search. Below, we plot the difference between the present and optimal function against the number of function evaluations required.



4. Conclusion

Randomization provides an interesting perspective for a variety of algorithms. Consider the perspective adhered to in this article in which our cost measure is the number of function evaluations required, assuming line searches are relatively cheap being a one-dimensional optimization problem, and with derivative information only available through (and well-approximated by) finite differencing. It was then shown in Section 1 that random search is comparable to steepest descent. Then, using the Hessian estimation technique introduced in Section 2, Section 3 demonstrated how these techniques can be used to accelerate random search. Finally, we devised a conjugate directions algorithm that incorporates second derivative information without directly requiring gradient information while sharing certain behaviours with more traditional Newton-like methods.

We make no claim that the conceptual techniques described earlier, in their pure form, are competitive with already-known derivative-based or derivative-free algorithms. We simply intend to illustrate how incorporating randomization provides a novel approach to the design of algorithms, even in very simple optimization schemes, suggesting that it may deserve further consideration. Note that all the algorithms considered in this paper, at each iteration, require only directional derivative or directional second-order information, creating a connection between the realms of derivative-free and derivative-based algorithms when this derivative information is well-approximated by finite differencing.

Acknowledgements

The authors would like to thank an anonymous referee for the helpful suggestions on substantially improving the presentation of this article. This research was supported in part by National Science Foundation Grant DMS-0806057.

References

- [1] H. Akaike, *On a successive transformation of probability distribution and its application to the analysis of the optimum gradient method*, Ann. Inst. Stat. Math. 11 (1959), pp. 1–16.
- [2] P. Billingsley, *Probability and Measure*, John Wiley & Sons, New York, 1986.
- [3] K.W. Brodlić, *A new direction set method for unconstrained minimization without evaluating derivatives*, J. Inst. Math. Appl. 15 (1975), pp. 385–396.
- [4] Yu. Ermoliev and R.J.-B. Wets, *Numerical Techniques for Stochastic Optimization*, Springer-Verlag, Berlin, 1988.
- [5] M. Gaviano, *Some general results on convergence of random search algorithms in minimization problems*, in *Towards Global Optimisation*, L.C.W. Dixon and G.P. Szego, eds., North-Holland, London, 1975, pp. 149–157.
- [6] Z.Q. Luo and P. Tseng, *On the convergence of the coordinate descent method for convex differentiable minimization*, J. Optim. Theory Appl. 72 (1992), pp. 7–35.
- [7] L. Nazareth, *Generation of conjugate directions for unconstrained minimization without derivatives*, Math. Comput. 30(133) (1976), pp. 115–131.
- [8] M.J.D. Powell, *An efficient method for finding the minimum of a function of several variables without calculating derivatives*, Comput. J. 7 (1964), pp. 155–162.
- [9] M.J.D. Powell, *On search directions for minimization algorithms*, Math. Progr. 4 (1973), pp. 193–201.
- [10] G. Rappl, *On linear convergence of a class of random search algorithms*, ZAMM – J. Appl. Math. Mech. 69(1) (1989), pp. 37–45.
- [11] F.J. Solis and R.J.-B. Wets, *Minimization by random search techniques*, Math. Oper. Res. 6(1) (1981), pp. 19–30.
- [12] J.C. Spall, *Introduction to Stochastic Search and Optimization: Estimation, Simulation and Control*, John Wiley & Sons, New York, 2003.
- [13] W.I. Zangwill, *Minimizing a function without calculating derivatives*, Comput. J. 10(3) (1967), pp. 293–296.

Appendix

The appendix to this article is dedicated to prove Theorem 3.2, the superlinear convergence of Algorithm 3.1 on convex quadratic functions. Before doing so, we define some notation. For a matrix, A , the *spectral norm* of A is the quantity $\|A\|_2 := \max_{\|x\|_2=1} \|Ax\|_2$ and the *Frobenius norm* of A is given by $\|A\|_F := \sqrt{\sum_{i,j} a_{ij}^2}$. If A is invertible, then $\|A^{-1}\|_2$ can also be expressed as the smallest constant K such that $\|Ax\|_2 \geq \frac{1}{K} \|x\|_2$ for all vectors x . Let $\text{Diag}(A)$ be the matrix whose diagonal matches that of A and whose non-diagonal entries are zero. Further, these norms satisfy

$$\|A\|_2 \leq \|A\|_F \leq \sqrt{n} \|A\|_2, \quad (\text{A.1})$$

$$\|Ax\|_2^2 \leq \lambda_{\max}(A) \|x\|_A^2 \leq \lambda_{\max}(A)^2 \|x\|_2^2 \quad (\text{A.2})$$

and

$$\lambda_{\min}(A) \|x\|_2^2 \leq \|x\|_A^2 \quad (\text{A.3})$$

where $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ are the maximum and minimum eigenvalues of A , respectively.

The proof of Theorem 3.2 will be broken into several lemmas, as follows.

LEMMA A.4 *Let A be positive definite and $\{v_1, \dots, v_n\}$ be unit norm vectors satisfying $|v_i^T A v_j| < \epsilon$ for all $i \neq j$. Then, for all $\epsilon > 0$ sufficiently small, the matrix $V = [v_1, v_2, \dots, v_n]$ is invertible and $\|V^{-1}\|$ is uniformly bounded above.*

Proof Consider any $y \in \mathbf{R}^n$ such that $\|y\| = 1$. Then,

$$\begin{aligned} \|Vy\|_2^2 &= y^T V^T V y \\ &= (y^T V^T A^{\frac{1}{2}}) A^{-1} (A^{\frac{1}{2}} V y) \\ &= \|A^{\frac{1}{2}} V^T y\|_{A^{-1}}^2 \\ &\geq \lambda_{\min}(A^{-1}) \|A^{\frac{1}{2}} V y\|^2 \quad (\text{by inequality (A.3)}) \\ &= \frac{1}{\lambda_{\max}(A)} y^T V^T A V y \\ &= \frac{1}{\lambda_{\max}(A)} [y^T \text{Diag}(V^T A V) y + y^T [V^T A V - \text{Diag}(V^T A V)] y] \\ &\geq \frac{\lambda_{\min}(A) - n\epsilon}{\lambda_{\max}(A)}, \end{aligned}$$

with the last inequality coming from the fact that

$$y^T \text{Diag}(V^T A V) y = \sum_{i=1}^n y_i^2 v_i^T A v_i,$$

inequality (A.3) and the fact that $|v_i^T A v_j| < \epsilon$ for all $i \neq j$. From the above bound and the alternative definition of $\|V^{-1}\|_2$, it follows that

$$\|V^{-1}\|_2^2 \leq \frac{\lambda_{\max}(A)}{\lambda_{\min}(A) - n\epsilon}.$$

In particular, for ϵ sufficiently small, the result follows. ■

LEMMA A.5 *Let A be positive definite and $\{v_1, \dots, v_n\}$ be unit norm vectors satisfying $|v_i^T A v_j| < \epsilon$ for all $i \neq j$. Then, for all $\epsilon > 0$ sufficiently small, there exists a matrix M such that $v_i^T M v_j = 0$ for all $i \neq j$, $v_i^T M v_i = v_i^T A v_i$ for all i and $\|A - M\| = O(\epsilon)$.*

Proof Let $V = [v_1, v_2, \dots, v_n]$ and recall that, by Lemma A.4, V is invertible for all $\epsilon > 0$ sufficiently small. Consider the matrix M defined by

$$M = A - V^{-T}(V^T A V - \text{Diag}(V^T A V))V^{-1} = V^{-T} \text{Diag}(V^T A V) V^{-1}.$$

It is clear by construction that $v_i^T M v_j = 0$ for all $i \neq j$ and that $v_i^T M v_i = v_i^T A v_i$. Further, observe that

$$\begin{aligned} \|A - M\|_2 &= \|V^{-T}(V^T A V - \text{Diag}(V^T A V))V^{-1}\|_2 \\ &\leq \|V^{-1}\|_2^2 \|(V^T A V - \text{Diag}(V^T A V))\|_F \\ &\leq \|V^{-1}\|_2^2 n\epsilon \end{aligned}$$

with the first inequality coming from the sub-multiplicity property for the spectral norm, the fact that the spectral norm is invariant under matrix transposition and inequality (A.1) while the last inequality comes from the fact that $|v_i^T A v_j| < \epsilon$ for all $i \neq j$. In particular, since $\|V^{-1}\|$ is bounded above for all $\epsilon > 0$ sufficiently small by Lemma A.4, the matrix M satisfies $\|A - M\|_2 = O(\epsilon)$ and the result follows. ■

Proof of Theorem 3.2 Suppose that, at iteration k of the algorithm, B_k is the current Hessian estimate. Define $\epsilon_k = \|B_k - A\|_F$ and note that, by inequality (A.1), $\|B_k - A\|_2 \leq \epsilon_k$.

Consider n consecutive iterations of the algorithm, beginning at iterate x_k and ending at iterate x_{k+n} . Without loss of generality, assume that the respective search directions satisfy $\|v_i\| = 1$ for $i = k, k+1, \dots, k+n-1$. Recall that by design of the algorithm, these search directions satisfy $v_i^T B_j v_j = 0$ for any $j \in \{k, k+1, \dots, k+n-1\}$ and $i \in \{j-n+1, \dots, j-1\}$. Note that this implies that for any $i < j \in \{k, k+1, \dots, k+n-1\}$,

$$|v_i^T A v_j| = |v_i^T B_j v_j + v_i^T (A - B_j) v_j| \leq \|v_i\| \|A - B_j\| \|v_j\| \leq \epsilon_k \quad (\text{A.4})$$

by inequality (A.1) and the definition of ϵ_k .

At each iteration $i = k, k+1, \dots, k+n-1$, Algorithm 3.1 obtains the new point by way of exact line search, getting

$$\begin{aligned} x_{i+1} &= x_i - \frac{v_i^T (A x_i + b)}{v_i^T A v_i} v_i \\ &= x_i - \frac{v_i^T (A x_k + b + \sum_{j=k}^{i-1} \alpha_j A v_j)}{v_i^T A v_i} v_i \\ &= x_i - \frac{v_i^T (A x_k + b)}{v_i^T A v_i} v_i - \frac{\sum_{j=k}^{i-1} \alpha_j v_i^T A v_j}{v_i^T A v_i} v_i, \end{aligned}$$

where α_j is defined by $\alpha_j = -\frac{v_j^T \nabla f(x_j)}{v_j^T A v_j}$. Expanding this over n consecutive iterations, we obtain

$$x_{k+n} = x_k - \sum_{i=k}^{k+n-1} \frac{v_i^T \nabla f(x_k)}{v_i^T A v_i} v_i - \sum_{i=k}^{k+n-1} \sum_{j=k}^{i-1} \frac{\alpha_j v_i^T A v_j}{v_i^T A v_i} v_i.$$

In particular, this implies

$$\begin{aligned} \|x_{k+n} - x^*\| &\leq \left\| x_k - \sum_{i=k}^{k+n-1} \frac{v_i^T \nabla f(x_k)}{v_i^T A v_i} v_i - x^* \right\| \\ &\quad + \sum_{i=k}^{k+n-1} \sum_{j=k}^{i-1} \left| \frac{v_j^T \nabla f(x_j)}{v_j^T A v_j} \frac{v_i^T A v_j}{v_i^T A v_i} \right|. \end{aligned} \quad (\text{A.5})$$

Let $V_k = [v_k, v_{k+1}, \dots, v_{k+n-1}]$. Since $B_k \rightarrow A$ almost surely by Corollary 2.2, it follows that for k sufficiently large, ϵ_k is sufficiently small almost surely. Therefore, suppose that k is sufficiently large so that V_k is invertible by Lemma A.4 and define M_k as in Lemma A.5. Recall that since v_k, \dots, v_{k+n-1} are conjugate with respect to M_k and $v_i^T A v_i = v_i^T M_k v_i$, it follows that

$$\left\| x_k - \sum_{i=k}^{k+n-1} \frac{v_i^T \nabla f(x_k)}{v_i^T A v_i} v_i - x^* \right\| = \|x_k - M_k^{-1} \nabla f(x_k) - x^*\|. \quad (\text{A.6})$$

Next, recall that since the algorithm is implemented with an exact line search, the objective function is non-increasing at each iteration. Specifically, for all j , $f(x_{j+1}) \leq f(x_j)$. By Equation (1.1), it can be seen that

$$\frac{1}{2} \|x_{j+1} - x^*\|_A^2 = f(x_{j+1}) - f(x^*) \leq f(x_j) - f(x^*) = \frac{1}{2} \|x_j - x^*\|_A^2.$$

This implies that the sequence $\{x_k\}_{k \geq 0}$ is bounded. Additionally, along with inequality (A.2), this implies that for $j \geq k$ we have,

$$\begin{aligned} |v_j^T \nabla f(x_j)| &\leq \|v_j\| \|\nabla f(x_j)\| \\ &= \|A(x_j - x^*)\| \\ &\leq \sqrt{\lambda_{\max}(A)} \|x_j - x^*\|_A \\ &\leq \sqrt{\lambda_{\max}(A)} \|x_k - x^*\|_A \\ &\leq \lambda_{\max}(A) \|x_k - x^*\|. \end{aligned}$$

Combining the above inequality, inequalities (A.3)–(A.6) and Equation (A.8), we conclude that

$$\|x_{k+n} - x^*\| \leq \|x_k - M_k^{-1} \nabla f(x_k) - x^*\| + \frac{n^2 \lambda_{\max}(A)}{\lambda_{\min}^2(A)} \|x_k - x^*\| \epsilon_k. \quad (\text{A.7})$$

Further, observe that since $\nabla f(x_k) = Ax_k + b = A(x_k - x^*)$, it follows that

$$\|x_k - M_k^{-1} \nabla f(x_k) - x^*\| = \|(I - M_k^{-1} A)(x_k - x^*)\|. \quad (\text{A.8})$$

From this, the fact that $\{x_k - x^*\}_{k \geq 0}$ is bounded, and the fact that $M_k \rightarrow A$ almost surely, it follows that

$$\|x_k - M_k^{-1} \nabla f(x_k) - x^*\| \rightarrow 0$$

almost surely. Combining this result, the fact that $\epsilon_k \rightarrow 0$ almost surely and inequality (A.7), it follows that $\|x_{k+n} - x^*\| \rightarrow 0$ almost surely, proving that the algorithm converges almost surely.

Finally, consider scaling inequality (A.7) by $\|x_k - x^*\|$, obtaining

$$\frac{\|x_{k+n} - x^*\|}{\|x_k - x^*\|} \leq \frac{\|x_k - M_k^{-1} \nabla f(x_k) - x^*\|}{\|x_k - x^*\|} + \frac{\frac{n^2 \lambda_{\max}(A)}{\lambda_{\min}^2(A)} \|x_k - x^*\| \epsilon_k}{\|x_k - x^*\|}.$$

Since, by Equation (A.8),

$$\frac{\|x_k - M_k^{-1} \nabla f(x_k) - x^*\|}{\|x_k - x^*\|} = \frac{\|(I - M_k^{-1} A)(x_k - x^*)\|}{\|x_k - x^*\|},$$

it follows that the first term converges to zero almost surely since $M_k \rightarrow A$ almost surely. Further, since $\epsilon_k \rightarrow 0$ almost surely, the second term converges to zero almost surely. These two facts together imply that

$$\frac{\|x_{k+n} - x^*\|}{\|x_k - x^*\|} \rightarrow 0$$

almost surely: by definition, this means the algorithm is n -step superlinearly convergent almost surely. ■