ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, Illinois 60439

Estimating Derivatives of Noisy Simulations[1]

Jorge J. Moré and Stefan M. Wild

Mathematics and Computer Science Division

Preprint ANL/MCS-P1785-0810

August 2010

[1]Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439. This work was supported by the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract DE-AC02-06CH11357.

Contents

Estimating Derivatives of Noisy Simulations[1]

Jorge J. Moré and Stefan M. Wild

Abstract

We employ recent work on computational noise to obtain near-optimal finite difference estimates of the derivatives of a noisy function. Our analysis employs a stochastic model of the noise without assuming a specific form of distribution. We use this model to derive theoretical bounds for the errors in the difference estimates and obtain an easily computable difference parameter that is provably near-optimal. Numerical results closely resemble the theory and show that we obtain accurate derivative estimates even when the noisy function is deterministic.

1 Introduction

We consider a fundamental problem in scientific computing: Given $f : \mathbb{R}^n \mapsto \mathbb{R}^m$, a base point x_0, and a direction $p \in \mathbb{R}^n$, compute an approximation to the directional derivative $f'(x_0; p)$. Of special importance is the case where the evaluation of f is the result of noisy simulations. In these simulations the value of f is known only within a given tolerance as a result of finite precision iterative methods and adaptive strategies, and this uncertainty in the value of f gives rise to computational noise. We are also interested in situations where the evaluation of f is computationally expensive.

There are basically three approaches to approximating directional derivatives. We can hand-code the derivative and thus be assured of a high-precision approximation. This approach is error-prone, however, and often considered infeasible for complex simulations. We can use automatic differentiation techniques [1, 3, 7] to compute an approximation. Many users consider this approach to be inapplicable to complex simulations; but in our experience, this view tends to be based on an incomplete understanding of automatic differentiation. If properly implemented, automatic differentiation techniques provide accurate approximations to the derivative. Automatic differentiation, however, can be computationally expensive relative to the evaluation of f. The relative cost depends on the application that produces f; but in theory, automatic differentiation techniques can produce a directional derivative at a modest multiple of the cost of evaluating f.

In this work we study how to obtain accurate difference approximations to the directional derivative. This approach is easily implemented and produces good approximations if the difference parameter is chosen appropriately; as we will demonstrate, the accuracy depends

[1]Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439. This work was supported by the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract DE-AC02-06CH11357.

on the noise level of the function. Another advantage is that a difference approximation is potentially less expensive than automatic differentiation techniques.

We motivate this work with the approximation of Jacobian-vector products for Newton-Krylov solvers, Given $f : \mathbb{R}^n \mapsto \mathbb{R}^n$ and a vector $p \in \mathbb{R}^n$, the approximation

$$\frac{f(x_0 + tp) - f(x_0)}{t}$$

can be used to estimate the Jacobian-vector product $f'(x_0)p$. In current practice, the choice of difference parameter is usually of the form

$$\sigma(x_0)\varepsilon_M^{1/2},$$

where ε_M is the unit roundoff and $\sigma(\cdot)$ is a scaling factor. A popular choice [9] for the scaling factor is

$$\sigma(x_0) = \frac{\max\left\{|x_0^T p|, x_*^T |p|\right\}}{\|p\|} \operatorname{sign}\left(x_0^T p\right),$$

where x_* is a vector of typical values for the absolute values of the solution. The survey by Knoll and Keyes [10] discusses other choices and notes that ε_M should be replaced by ε_{rel} when f can be evaluated only to a relative precision of ε_{rel}, but the authors do not indicate how to determine ε_{rel}.

Our aim in this work is to provide a precise specification for a difference parameter that is computationally feasible and provably nearly optimal. Our approach bridges early work [12, 6, 5] that assumes that the computed function f is deterministic with recent work [2] where f is assumed to be determined by a stochastic process. We emphasize the situation where the computed f is the result of a deterministic noisy simulation since this case invariably arises in computationally expensive, complex simulations.

Directional derivatives can be studied by restricting attention to scalar-valued functions defined in an interval I; the vector-valued case where $f : \mathbb{R} \mapsto \mathbb{R}^m$ is handled by working with each component of f. Our computational model assumes that the computed function f can be expressed as

$$f(t) = f_s(t) + \varepsilon(t), \qquad t \in I, \tag{1}$$

where $f_s : \mathbb{R} \mapsto \mathbb{R}$ is a smooth, deterministic function and the noise $\varepsilon : \mathbb{R} \mapsto \mathbb{R}$ is stochastic. The random variables $\varepsilon(t)$ are assumed to be independent and identically distributed (iid) for all $t \in I$, and we define the noise level of the function f as the standard deviation

$$\varepsilon_f = (\operatorname{Var}\{\varepsilon(t)\})^{1/2} \tag{2}$$

of the noise. We used this model [13] to study computational noise in both stochastic and deterministic simulations. As part of this study we developed the ECnoise algorithm to determine the noise level ε_f in a few function evaluations. We provide additional information on computational noise in Section 2.

We formulate an approximation problem in terms of the derivative of the expected value $\operatorname{E}\{f\}$ of the computed function f. Our assumptions on the model (1) imply that

$\mathrm{E}\{f\} = f_s + \mu$ where μ is the mean of the noise, and thus the derivatives of f_s and $\mathrm{E}\{f\}$ agree. Hence, we approximate the derivative of the expected value by choosing a difference parameter h that minimizes the (squared) l_2 error

$$\mathcal{E}(h) = \left(\frac{f(t_0 + h) - f(t_0)}{h} - f_s'(t_0) \right)^2$$

for all h such that $t_0 + h \in I$. Minimizing the expected value $\mathrm{E}\{\mathcal{E}(h)\}$ yields an optimal approximation to the derivative of the expected value $\mathrm{E}\{f\}$. Our results will show that we can obtain nearly optimal estimates of the derivative from rough estimates of the noise level and $|f''|$.

Section 3 presents the main results for estimating the derivative. We show that the minimal value of the expected value $\mathrm{E}\{\mathcal{E}(h)\}$ lies in the interval $\gamma_1 \varepsilon_f [\mu_L, \mu_M]$, where γ_1 is a constant and (μ_L, μ_M) are the minimum and maximum of $|f''|$ on I, respectively. This result provides tight bounds on the best possible error, and thus we use the term *nearly optimal* for any difference parameter h such that $\mathrm{E}\{\mathcal{E}(h)\}$ lies in this interval.

As a consequence of the results in Section 3 we show that the expected best possible l_2 error in a forward difference approximation to f' is of order $\varepsilon_f^{1/2}$ and that a difference parameter h^* achieves this error. This result provides further justification for the use of the term *nearly optimal*. In the remainder of this section we study, in particular, the variance $\mathrm{Var}\{\mathcal{E}(h)\}$ since this variance characterizes the spread of the errors about the mean.

The results in Section 3 for the forward difference approximation of f' can be extended to other approximations of f' or to approximations of higher-order derivatives. Section 4 illustrates these extensions with the central difference approximation to f' and f''. In both cases we determine the best possible l_2 error in the difference approximation.

Section 5 describes our algorithm for obtaining a nearly optimal estimate of the derivative in a deterministic simulation. We consider forward differences but a similar approach can be applied for other difference schemes. The main ingredient in this algorithm is a test to decide whether a difference parameter h produces an adequate estimate of f''. Given this estimate and an estimate for the noise level ε_f, we determine an estimate h^* for the optimal difference parameter.

Section 6 presents computational experiments for both stochastic and deterministic problems. The aim of these experiments is to study the performance of the parameter h^* determined in Section 5. Our sample problems include a Monte Carlo simulation, noisy quadratic problems defined by the iterative (bicgstab) solution of sparse systems of linear equations, and smooth nonlinear problems where the number of variables range up to $6.4 \cdot 10^5$. In all cases we show that $\mathcal{E}(h^*) \le \mathcal{E}(h)$ for almost every h.

Our computational results show that we can produce nearly optimal approximations if we have rough estimates of the noise level and $|f''|$. We claim that once these estimates are obtained at a base point x_0, they will be valid in a reasonably large neighborhood of x_0 so that new estimates will be needed only if the underlying algorithm makes large changes to x_0. Integration of our results into algorithms is a research topic that we plan to explore.

2 Background

We review the theoretical and computational framework [13] for the study of computational noise in a neighborhood $N(x_0)$ of a base point x_0. The main assumption is that

$$f(x) = f_s(x) + \varepsilon(x), \qquad x \in N(x_0), \tag{3}$$

where $f : \mathbb{R}^n \mapsto \mathbb{R}$ is the computed function, $f_s : \mathbb{R}^n \mapsto \mathbb{R}$ is a smooth deterministic function, and the noise $\varepsilon : \mathbb{R}^n \mapsto \mathbb{R}$ is a random variable whose distribution is independent of x. The standard deviation $\varepsilon_f = (\mathrm{Var}\{\varepsilon(x)\})^{1/2}$ is then the *noise level* of the function. This model of computational noise assumes that f is a stochastic process where the output of the simulation is a (random) variable. As we will see, this model can provide useful results even when f is deterministic.

The noise level of a function f provides the standard deviation for the values of a simulation defined by f. This interpretation of the noise level has a rigorous justification if the output $f(x)$ of a simulation is a random variable with an expected value of $\mathrm{E}\{f(x)\}$ and standard deviation ε_f. In this case the Chebyshev inequality

$$\mathcal{P}\big\{|f(x) - \mathrm{E}\{f(x)\}| \le \gamma\varepsilon_f\big\} \ge 1 - \frac{1}{\gamma^2},$$

where $\mathcal{P}\{\cdot\}$ is the probability of the event, implies that

$$|f(x) - \mathrm{E}\{f(x)\}| \le \gamma\varepsilon_f \tag{4}$$

is likely to hold for $\gamma \ge 1$ of modest size. Thus, (4) holds in at least 99% of the cases with $\gamma = 10$. Of course, tighter bounds are available if we have additional information on the distribution of f. For example, if the distribution is normal, then (4) holds in at least 99.7% of the cases with $\gamma = 3$.

We have developed the ECnoise algorithm [13] to determine the noise level ε_f of a function. The theoretical framework of ECnoise is based on stochastic noise but, importantly, does not assume a specific distribution for the noise. On deterministic simulations our computational results show that ECnoise produces reliable results in few function evaluations.

While the noise level ε_f is not a bound on rounding errors, it can be related to an absolute bound ε_A if one is willing to make further distributional assumptions. For example, if the techniques for estimating ε_A in [5, 6, 12] are interpreted in a stochastic framework by assuming that the values of $f(t)$ are uniformly distributed in $[f(t) - \varepsilon_A, f(t) + \varepsilon_A]$, then the absolute error may be obtained from the noise level by

$$\varepsilon_A = (3^{1/2})\varepsilon_f.$$

The absolute error can be related to the noise level in a similar way for other distributions with compact support.

To illustrate some of the features of computational noise, we consider the function defined in Figure 1. This deterministic function analytically computes the square of a real

```
f = t;
for k = 1:L; f = sqrt(f); end
for k = 1:L; f = f^2;     end
f = f^2;
```

Figure 1: Code for the **higham** function $f(t)$.

number in a manner inspired by Higham [8, pages 15–16]. As the parameter L increases, the noise in f tends to increase.

Figure 2 (left) shows the error $f(t) - t^2$ for $t \approx 2$ when f is computed in MATLAB's double precision with an AMD 64-bit processor and $L = 30$. We note that f systematically underestimates t^2 and hence these errors should not be considered to have zero mean. Furthermore, though the errors appear to be relatively uniform, they are clearly not independent and identically distributed in space. This is the central assumption in both the present work and in [13]; but, as in [13], we will show that we obtain reasonable solutions even when it is violated. For this example, the ECnoise algorithm estimates the noise level to be $\varepsilon_f = 4.9 \cdot 10^{-7}$. As discussed earlier, ε_f is not a bound on the errors, and Figure 2 (left) shows errors as large as $18 \cdot 10^{-7}$.

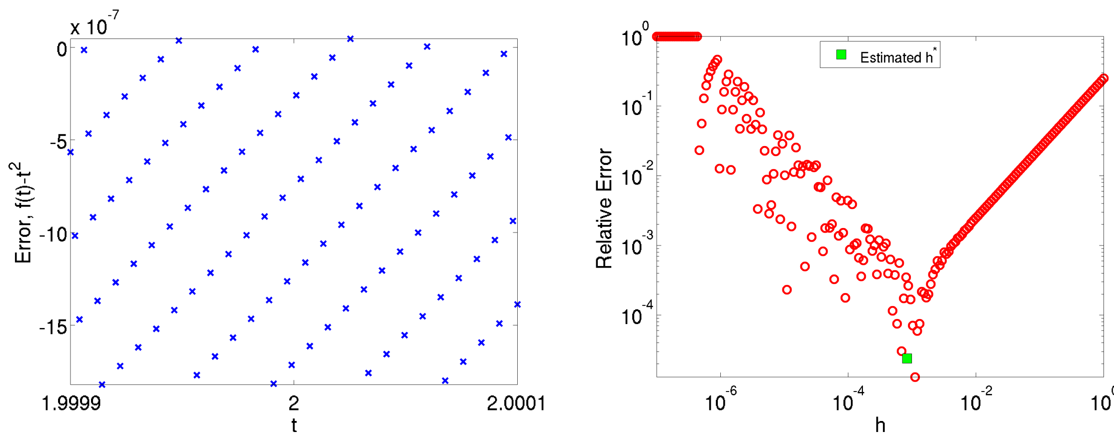

Figure 2: The noise in f (left) and the relative error in forward difference estimates of the derivative for a variety of difference sizes h (right).

Figure 2 (right) is the subject of the present work and shows the relative errors between a forward difference estimate for $f'(t)$ at $t = 2$, for difference parameters ranging from $h = 10^{-7}$ to $h = 1$, when compared with the analytic derivative $2t$. The behavior shown is typical for noisy functions. When h is large, the higher-order derivatives can lead to larger errors; while for smaller h, the derivative estimates can vary greatly because the differences in the function are dominated by the noise. For sufficiently small values of h, which in this case includes the commonly used square root of machine precision $h = 1.5 \cdot 10^{-8}$, we compute $f(t + h) = f(t)$. Thus the estimate of $f'(t)$ is 0.

Figure 2 (right) illustrates that a judicious choice of h can lead to a derivative estimate expected to be correct to four digits. Our goal in this work is to obtain such an h estimate with only a few additional function values.

3 Estimating Directional Derivatives

Our results for estimating the derivative of a computed function $f : \mathbb{R} \mapsto \mathbb{R}$ use the computational model (1) for f so that the noise level of f is defined by (2). We implicitly assume that the variance $\mathrm{Var}\{\varepsilon(t)\}$ is finite so that estimation of ε_f is reasonable and that $\varepsilon_f > 0$.

We first consider forward difference approximations to the derivative; in the next section we show that the ideas extend to centered and higher-order approximations. For forward difference approximations, we assume that the interval I in (1) is of the form

$$I = \left\{ t : t \in [t_0, t_0 + h_0] \right\}$$

for some t_0 and $h_0 > 0$, and we seek a difference parameter $h > 0$ that minimizes the expected value $\mathrm{E}\{\mathcal{E}(h)\}$ of the least squares error

$$\mathcal{E}(h) = \left(\frac{f(t_0 + h) - f(t_0)}{h} - f_s'(t_0) \right)^2 . \tag{5}$$

As noted in the introduction, this criterion requires that we approximate the derivative of the expected value $\mathrm{E}\{f\}$ at t_0. We will show that it is an excellent predictor of computational performance.

We need some basic properties of the expectation and variance operators. Recall that the expectation is a linear operator and that for any scalars α_k

$$\mathrm{Var}\left\{ \sum_{k=1}^{m} \alpha_k X_k \right\} = \sum_{k=1}^{m} \alpha_k^2 \mathrm{Var}\{X_k\}$$

if the random variables $X_1, \ldots, X_m$ are independent. Also, for a random variable X with mean $\mathrm{E}\{X\} = 0$, we have

$$\mathrm{E}\left\{ (X + \alpha)^2 \right\} = \mathrm{Var}\{X\} + \alpha^2 \tag{6}$$

for any constant α.

Lemma 3.1 *Assume that $f_s : \mathbb{R} \mapsto \mathbb{R}$ is twice differentiable in I, and set μ_L and μ_M to the minimum and maximum of $|f_s''|$ on I, respectively. If $h \in (0, h_0]$ then*

$$\phi(h; \mu_L) \le \mathrm{E}\{\mathcal{E}(h)\} \le \phi(h; \mu_M),$$

where

$$\phi(h; \mu) = \frac{1}{4}\mu^2 h^2 + 2\frac{\varepsilon_f^2}{h^2}.$$

Proof. Taylor's theorem shows that

$$\frac{f_s(t_0+h)-f_s(t_0)}{h} = f_s'(t_0) + \tfrac{1}{2} f_s''(\xi) h$$

for some $\xi \in [t_0, t_0+h] \subseteq I$, and thus

$$\mathcal{E}(h) = \left(\tfrac{1}{2} f_s''(\xi) h + \frac{\varepsilon(t_0+h)-\varepsilon(t_0)}{h} \right)^2 .$$

Since $h \le h_0$, the iid assumption on the noise $\varepsilon(\cdot)$ and (6) then yield

$$\mathrm{E}\left\{\mathcal{E}(h)\right\} = \frac{1}{4} f_s''(\xi)^2 h^2 + 2\frac{\varepsilon_f^2}{h^2}.$$

The expression on the right cannot be optimized easily because ξ depends on h, but we can estimate this expression in terms of μ_L and μ_M. Indeed, the expression implies that

$$\frac{1}{4}\mu_L^2 h^2 + 2\frac{\varepsilon_f^2}{h^2} \le \mathrm{E}\left\{\mathcal{E}(h)\right\} \le \frac{1}{4}\mu_M^2 h^2 + 2\frac{\varepsilon_f^2}{h^2}.$$

This is the desired result. ■

Lemma 3.1 provides the background needed to obtain sharp estimates of the minimum of the expectation of the least squares deviation defined by (5). A key observation is that ϕ is uniformly convex and that

$$\min_h \phi(h;\mu) = \gamma_1 \mu \varepsilon_f, \qquad \gamma_1 = 2^{1/2}. \tag{7}$$

In addition to $\varepsilon_f > 0$, we assume that $\mu_M > 0$ since otherwise f_s is affine and setting $h = h_0$ would suffice. For future reference note that the global minimizer of ϕ for $\mu = \mu_M$ is

$$h_M = \gamma_2 \left(\frac{\varepsilon_f}{\mu_M} \right)^{1/2}, \qquad \gamma_2 = 8^{1/4} \approx 1.68. \tag{8}$$

The following result shows that h_M plays an important role in the analysis of $h \mapsto \mathrm{E}\left\{\mathcal{E}(h)\right\}$.

Theorem 3.2 *Assume that $f_s : \mathbb{R} \mapsto \mathbb{R}$ is twice differentiable in I; set μ_L and μ_M to the minimum and maximum of $|f_s''|$ on I, respectively; and define h_M by* (8)*. If $h_M \le h_0$, then*

$$\gamma_1 \mu_L \varepsilon_f \le \min_{0 \le h \le h_0} \mathrm{E}\left\{\mathcal{E}(h)\right\} \le \gamma_1 \mu_M \varepsilon_f.$$

Proof. Lemma 3.1 implies that $\phi(h;\mu_L) \le \mathrm{E}\left\{\mathcal{E}(h)\right\} \le \phi(h;\mu_M)$ for $h \le h_0$, and thus

$$\min_{0 \le h \le h_0} \phi(h;\mu_L) \le \min_{0 \le h \le h_0} \mathrm{E}\left\{\mathcal{E}(h)\right\} \le \min_{0 \le h \le h_0} \phi(h;\mu_M).$$

We estimate the lower bound by using (7) to show that

$$\gamma_1 \mu_L \varepsilon_f \le \min_{0 \le h \le h_0} \phi(h;\mu_L).$$

Moreover, since $h_M \le h_0$,

$$\min_{0 \le h \le h_0} \phi(h;\mu_M) = \gamma_1 \mu_M \varepsilon_f.$$

The result is now follows directly from the last two estimates. ■

An important consequence of Theorem 3.2 is that we can expect to obtain only an approximation to $f_s'(t_0)$ of order $\varepsilon_f^{1/2}$. This follows from Theorem 3.2 since

$$(\gamma_1 \mu_L \, \varepsilon_f)^{1/2} \le \min_{0 \le h \le h_0} \mathrm{E}\left\{\mathcal{E}(h)\right\}^{1/2} = \left(\min_{0 \le h \le h_0} \mathrm{E}\left\{\mathcal{E}(h)\right\}\right)^{1/2} \le (\gamma_1 \mu_M \, \varepsilon_f)^{1/2}.$$

Our numerical results will confirm this conclusion.

We claim that the assumption that $h_M \le h_0$ in Theorem 3.2 is not restrictive, but the main argument to support this claim is based on computational experiments. In practice we compute h_M by estimating the noise ε_f and approximating μ_M by an estimate of $|f_s''(t_0)|$. Thus, we obtain an estimate of h_M from (8). In our experiments we have found that, remarkably,

$$\min_{h \ge 0} \mathrm{E}\left\{\mathcal{E}(h)\right\} \approx \mathrm{E}\left\{\mathcal{E}(h_M)\right\},$$

so this supports the claim that $h_M \le h_0$ holds in practice.

We conclude this discussion of Theorem 3.2 with the observation that if $h_0 < h_M$. then the best possible error in $\mathcal{E}$ increases. This can be seen by noting that since $h_0 < h_L = \gamma_2 \sqrt{\varepsilon_f / \mu_L}$ in this case, h_0 is the solution of $\min_{0 \le h \le h_0} \phi(h;\mu_L)$, and thus Lemma 3.1 and (7) imply that

$$\gamma_1 \mu_L \varepsilon_f < \phi(h_0;\mu_L) \le \min_{0 \le h \le h_0} \phi(h;\mu_L) \le \min_{0 \le h \le h_0} \mathrm{E}\left\{\mathcal{E}(h)\right\}.$$

Thus, the lower bound in Theorem 3.2 increases for $h_0 < h_M$.

Theorem 3.2 provides bounds on the best possible error but does not provide a difference parameter h that minimizes $\mathrm{E}\left\{\mathcal{E}(h)\right\}$. We now show that h_M is an approximate minimizer in the sense that $\mathrm{E}\left\{\mathcal{E}(h_M)\right\}$ satisfies the bounds in Theorem 3.2.

Corollary 3.3 *Assume that $f_s : \mathbb{R} \mapsto \mathbb{R}$ is twice differentiable in I; set μ_L and μ_M to the minimum and maximum of $|f_s''|$ on I, respectively; and define h_M by* (8). *If $\mu_L > 0$ and $h_M \le h_0$, then*

$$\gamma_1 \mu_L \varepsilon_f \le \mathrm{E}\left\{\mathcal{E}(h_M)\right\} \le \gamma_1 \mu_M \varepsilon_f.$$

Moreover,

$$\mathrm{E}\left\{\mathcal{E}(h_M)\right\} \le \left(\frac{\mu_M}{\mu_L}\right) \min_{0 \le h \le h_0} \mathrm{E}\left\{\mathcal{E}(h)\right\}.$$

Proof. The bounds in Lemma 3.1 and (7) for μ_M and μ_L imply that

$$\gamma_1 \mu_L \, \varepsilon_f \le \phi(h_M;\mu_L) \le \mathrm{E}\left\{\mathcal{E}(h_M)\right\} \le \phi(h_M;\mu_M) = \gamma_1 \mu_M \, \varepsilon_f.$$

This proves the first claim in this result. We conclude the proof by noting that

$$\mathrm{E}\left\{\mathcal{E}(h_M)\right\} \le \gamma_1 \mu_M \, \varepsilon_f = \left(\frac{\mu_M}{\mu_L}\right) \gamma_1 \mu_L \, \varepsilon_f \le \left(\frac{\mu_M}{\mu_L}\right) \min_{0 \le h \le h_0} \mathrm{E}\left\{\mathcal{E}(h)\right\},$$

as desired. ■

Theorem 3.2 and Corollary 3.3 are new, but related results have appeared in the literature. Earlier work on estimating directional derivatives includes [5, 6, 12]. In these results it is assumed that the computed function is deterministic and of the form

$$f(t) = f_s(t) + e(t), \tag{9}$$

where the error function e represents the rounding errors in computing f_s in working precision. The main assumption in these results is that there is a uniform bound

$$|e(t)| \le \varepsilon_A, \qquad t \in I, \tag{10}$$

on the rounding errors. This assumption yields the bound

$$\left| \frac{f(t_0+h) - f(t_0)}{h} - f_s'(t_0) \right| \le \left| \tfrac{1}{2} f_s''(\xi) \right| |h| + 2\frac{\varepsilon_A}{|h|} \le \tfrac{1}{2}\mu_M |h| + 2\frac{\varepsilon_A}{|h|},$$

and thus the $h > 0$ that minimizes this upper bound on the l_1 error is

$$h_A = 2 \left(\frac{\varepsilon_A}{\mu_M} \right)^{1/2} .$$

In this approach the optimal parameter h_A minimizes an upper bound on the l_1 error, but we cannot rule out the possibility that there is another $h > 0$ that provides a significantly better reduction in the error. Also note that in this derivation it is implicitly assumed that $t_0 + h_A \in I$, that is, $h_A \le h_0$ when $I = [t_0, t_0 + h_0]$. Another theoretical difference is that ε_A is an absolute bound, and thus we may have $\varepsilon_A \gg \varepsilon_f$.

The recent contribution of Brekelmans et al. [2] assumes that the computed function is stochastic and of the form (1) where the noise is iid with mean zero on $\mathbb{R}$. This is a strong assumption in numerical simulations where the iid assumption is likely to hold only in a neighborhood of t_0. Several schemes for obtaining gradient estimates are analyzed [2] under this assumption; in the case of forward differences the upper bound

$$\mathrm{E}\left\{ \mathcal{E}(h) \right\} \le \phi(h; \mu_M)$$

is derived, where ϕ is defined as in Lemma 3.1, and then the optimal parameter is defined as the minimizer of the upper bound. This approach yields h_M as the optimal parameter, but there is no claim of optimality in the sense of Corollary 3.3. Moreover, since the assumptions on the noise term are global, these results assume that $h_0 = +\infty$.

Thus far we have pursued an h that minimizes the mean squared error without considering how bad a particular realization of the error for this optimal h could be. The following theorem characterizes the spread of the squared errors about this mean.

Theorem 3.4 *Assume that $f_s : \mathbb{R} \mapsto \mathbb{R}$ is twice differentiable in I; set μ_L and μ_M to the minimum and maximum of $|f_s''|$ on I, respectively; and define*

$$m_4 = \mathrm{E}\left\{ \left(\varepsilon(t) - \mathrm{E}\left\{ \varepsilon(t) \right\} \right)^4 \right\}$$

to be the fourth central moment of the noise. If $m_4 < \infty$, then for all $h \in (0, h_0]$,

$$2\mu_L^2 \varepsilon_f^2 + \frac{\gamma_3}{h^4} \le \operatorname{Var}\{\mathcal{E}(h)\} \le 2\mu_M^2 \varepsilon_f^2 + \frac{\gamma_3}{h^4},$$

where $\gamma_3 = 2(m_4 + \varepsilon_f^4)$.

Proof. We estimate $\operatorname{Var}\{\mathcal{E}(h)\}$ by computing $\operatorname{E}\{\mathcal{E}(h)\}^2$ and $\operatorname{E}\left\{\mathcal{E}(h)^2\right\}$. As in Lemma 3.1,

$$\mathcal{E}(h) = \left(\tfrac{1}{2} f_s''(\xi) h + \frac{\varepsilon(t_0 + h) - \varepsilon(t_0)}{h} \right)^2$$

for some $\xi \in [t_0, t_0 + h]$, and hence $h \in (0, h_0]$ implies that

$$\operatorname{E}\{\mathcal{E}(h)\}^2 = \frac{h^4}{16} f_s''(\xi)^4 + f_s''(\xi)^2 \varepsilon_f^2 + \frac{4\varepsilon_f^4}{h^4}.$$

Since $\varepsilon(t_0)$ and $\varepsilon(t_0 + h)$ are iid, a computation shows that

$$\operatorname{E}\{\varepsilon(t_0 + h) - \varepsilon(t_0)\} = 0, \qquad \operatorname{E}\left\{[\varepsilon(t_0 + h) - \varepsilon(t_0)]^2\right\} = 2\varepsilon_f^2,$$
$$\operatorname{E}\left\{[\varepsilon(t_0 + h) - \varepsilon(t_0)]^3\right\} = 0, \qquad \operatorname{E}\left\{[\varepsilon(t_0 + h) - \varepsilon(t_0)]^4\right\} = 2m_4 + 6\varepsilon_f^4.$$

These results, together with the above expression for $\mathcal{E}(h)$, yield that

$$\operatorname{E}\left\{\mathcal{E}(h)^2\right\} = \frac{h^4}{16} f_s''(\xi)^4 + 3 f_s''(\xi)^2 \varepsilon_f^2 + \frac{2}{h^4}\left(m_4 + 3\varepsilon_f^4\right).$$

Combining the expressions for $\operatorname{E}\{\mathcal{E}(h)\}^2$ and $\operatorname{E}\left\{\mathcal{E}(h)^2\right\}$, we obtain

$$\operatorname{Var}\{\mathcal{E}(h)\} = \operatorname{E}\left\{\mathcal{E}(h)^2\right\} - \operatorname{E}\{\mathcal{E}(h)\}^2 = 2 f_s''(\xi)^2 \varepsilon_f^2 + \frac{\gamma_3}{h^4}.$$

The bounds for $\operatorname{Var}\{\mathcal{E}(h)\}$ now directly follow from the definition of μ_L and μ_M as the extrema of $|f_s''|$ over the larger interval I within which ξ lies. ■

Theorem 3.4 extends a similar result in [2] by providing a lower bound for $\operatorname{Var}\{\mathcal{E}(h)\}$. Since $\gamma_3 > 0$, the bounds for $\operatorname{Var}\{\mathcal{E}(h)\}$ are both decreasing in h, and hence the variance is minimized on I at $h = h_0$. Thus, this result justifies choosing overestimates of h_M, as we will use in Section 5, instead of underestimates because of the tightening of the upper bound on the variance $\operatorname{Var}\{\mathcal{E}(h)\}$ as h increases. Also of interest are the bounds

$$\frac{1}{4}\mu_L^2 \left(9\varepsilon_f^2 + \frac{m_4}{\varepsilon_f^2} \right) \le \operatorname{Var}\{\mathcal{E}(h_M)\} \le \frac{1}{4}\mu_M^2 \left(9\varepsilon_f^2 + \frac{m_4}{\varepsilon_f^2} \right)$$

on the variance for h_M, which follow immediately from Theorem 3.4.

Before pursuing extensions to other difference estimates, we illustrate the bounds in the preceding results on a pair of simple stochastic functions,

$$f_2(t) = t^2 + 10^{-6} U_{[-2\sqrt{3}, 2\sqrt{3}]}, \qquad f_3(t) = t^3 + 10^{-6} U_{[-2\sqrt{3}, 2\sqrt{3}]}, \tag{11}$$

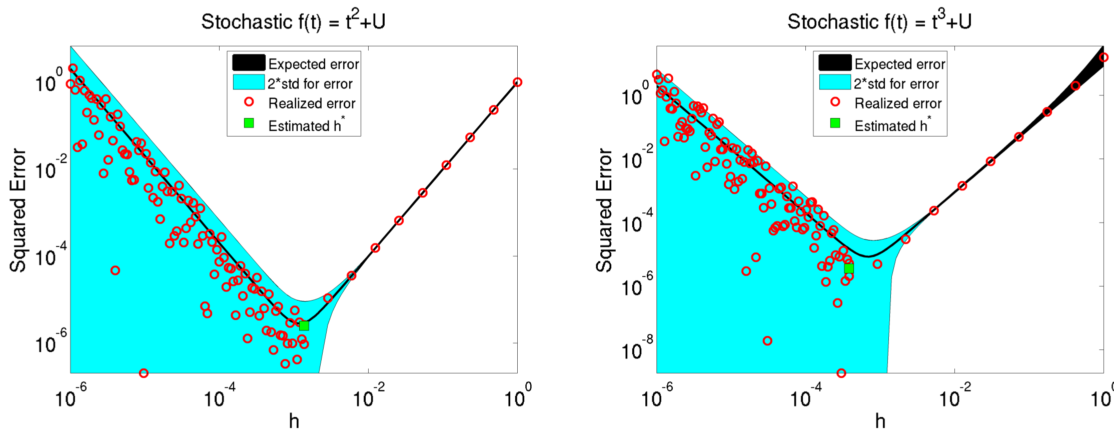

Figure 3: Log-log plots of realizations of $\mathcal{E}(h)$ for the stochastic functions f_2 and f_3 in (11) along with the expected error and uncertainty regions predicted by the theory.

where $U_{[a,b]}$ indicates a random variable distributed uniformly on $[a, b]$ and the interval $[a, b]$ was chosen so that $\varepsilon_f = 10^{-6}$ and $m_4 = 1.8 \cdot 10^{-24}$. We take $t_0 = 1$.

The circles in Figure 3 represent realizations of the (stochastic) error $\mathcal{E}(h)$ for different values of h and illustrate the trend and variability of the squared error as a function of h. Bounds, similar to those from Lemma 3.1,

$$\left[\phi\left(h; \min_{t_0 \le t \le t_0+h} |f_s''(t)|\right), \phi\left(h; \max_{t_0 \le t \le t_0+h} |f_s''(t)|\right)\right],$$

on the mean $\mathrm{E}\{\mathcal{E}(h)\}$ are shown in black. For the quadratic f_2, the second derivative f_s'' is constant, and hence the bounds are equal. For the cubic f_3, the rightmost part of Figure 3 shows that the width of the bounds grows as the difference between the minimum and maximum of $|f_s''|$ grows.

The shaded region in Figure 3 represents two standard deviations from the bounds on the mean, where the standard deviation is estimated by using a bound similar to that in Theorem 3.4,

$$\left(2\varepsilon_f^2 \max_{t_0 \le t \le t_0+h} |f_s''(t)|^2 + \frac{\gamma_3}{h^4}\right)^{1/2},$$

and shows that this variance becomes negligible relative to the mean $\mathrm{E}\{\mathcal{E}(h)\}$ as h grows. We also illustrate a realization of $\mathcal{E}(h^*)$ when h^* is defined by (8) using estimates of μ_M and ε_f described in Section 5.

4 Extensions

Our results for the forward difference approximation of f' can be extended to other approximations of f' or to approximations of higher-order derivatives. In this section we illustrate these extensions with the central difference approximation to f' and f''.

Our assumptions on the computed function are similar to those in Section 3 with the interval within which we assume that the random variables $\{\varepsilon(t) : t \in I\}$ are iid of the form

$$I = \left\{t : |t - t_0| \le h_0\right\}$$

for some t_0 and $h_0 > 0$. The techniques used in these extensions are similar to those used in Section 3. In all cases we obtain bounds of the form

$$\phi(h; \mu_L) \le \mathrm{E}\left\{\mathcal{E}(h)\right\} \le \phi(h; \mu_M),$$

where $\mathcal{E}$ is the squared error between the approximation and the derivative, and ϕ is a function of h and μ. We estimate the least value of $\mathrm{E}\left\{\mathcal{E}(h)\right\}$ from the two-sided bound of $\mathrm{E}\left\{\mathcal{E}(h)\right\}$ in terms of ϕ.

We first consider central difference approximations to f'. In this case the expected least squares error in the approximation is

$$\mathcal{E}_c(h) = \left(\frac{f(t_0 + h) - f(t_0 - h)}{2h} - f_s'(t_0)\right)^2.$$

Lemma 4.1 *Assume that $f_s : \mathbb{R} \mapsto \mathbb{R}$ is three times differentiable in I, and set μ_L and μ_M to the minimum and maximum of $|f_s^{(3)}|$ on I, respectively. If $|h| \le h_0$, then*

$$\phi_c(h; \mu_L) \le \mathrm{E}\left\{\mathcal{E}_c(h)\right\} \le \phi_c(h; \mu_M),$$

where

$$\phi_c(h; \mu) = \frac{1}{36}\mu^2 h^4 + \frac{1}{2}\frac{\epsilon_f^2}{h^2}.$$

Proof. A standard calculation shows that

$$\frac{f_s(t_0 + h) - f_s(t_0 - h)}{2h} = f_s'(t_0) + \frac{1}{6} f_s^{(3)}(\xi) h^2,$$

for some $\xi \in [t_0 - h, t_0 + h] \subseteq I$, and thus

$$\frac{f(t_0 + h) - f(t_0 - h)}{2h} = f_s'(t_0) + \frac{1}{6} f_s^{(3)}(\xi) h^2 + \frac{\varepsilon(t_0 + h) - \varepsilon(t_0 - h)}{2h}.$$

Since $|h| \le h_0$, the iid assumption on the noise $\varepsilon(\cdot)$ and (6) now yield that

$$\mathrm{E}\left\{\mathcal{E}_c(h)\right\} = \mathrm{E}\left\{\left(\frac{f(t_0 + h) - f(t_0 - h)}{2h} - f_s'(t_0)\right)^2\right\} = \left(\frac{1}{6} f_s^{(3)}(\xi)\right)^2 h^4 + \frac{1}{2}\frac{\epsilon_f^2}{h^2}.$$

The result follows by using μ_L and μ_M to bound the expression for $\mathrm{E}\left\{\mathcal{E}_c(h)\right\}$. ■

We estimate the minimum of the expected error $\mathrm{E}\left\{\mathcal{E}_c\right\}$ by analyzing the mapping ϕ_c in Lemma 4.1. A calculation shows that ϕ_c is uniformly convex and that

$$\min_h \phi_c(h; \mu) = \gamma_4 \mu^{2/3}\, \varepsilon_f^{4/3}, \qquad \gamma_4 = \frac{1}{4}\, 3^{1/3} \approx 0.361. \tag{12}$$

The next result shows that the global minimizer for μ_M, defined by

$$|h_M| = \gamma_5 \left(\frac{\varepsilon_f}{\mu_M}\right)^{1/3}, \qquad \gamma_5 = 3^{1/3} \approx 1.44, \tag{13}$$

plays an important role in the behavior of $\mathrm{E}\left\{\mathcal{E}_c\right\}$.

Theorem 4.2 *Assume that $f_s : \mathbb{R} \mapsto \mathbb{R}$ is three times differentiable in I; set μ_L and μ_M to the minimum and maximum of $|f_s^{(3)}|$ on I, respectively; and define h_M by (13). If $|h_M| \le h_0$, then*

$$\gamma_4\, \mu_L^{2/3}\, \varepsilon_f^{4/3} \le \min_{|h|\le h_0} \mathrm{E}\{\mathcal{E}_c(h)\} \le \gamma_4\, \mu_M^{2/3}\, \varepsilon_f^{4/3}.$$

Proof. Since $\gamma_4\, \mu^{2/3} \varepsilon_f^{4/3}$ is a lower bound for $\phi_c(\cdot, \mu)$, Lemma 4.1 shows that

$$\gamma_4\, \mu_L^{2/3} \varepsilon_f^{4/3} \le \min_{|h|\le h_0} \phi_c(h; \mu_L) \le \min_{|h|\le h_0} \mathrm{E}\{\mathcal{E}_c(h)\} \le \min_{|h|\le h_0} \phi_c(h; \mu_M).$$

The proof is completed by noting that (12) holds for $\mu = \mu_M$ when $|h_M| \le h_0$. ■

Theorem 4.2 shows that with a central difference approximation we can expect an error of order $\varepsilon_f^{2/3}$. This is an improvement over the order $\varepsilon_f^{1/2}$ error for one-sided differences that comes at a cost of one additional function evaluation. If f is noisy, this is not a significant improvement. For example, if $\varepsilon_f = 10^{-6}$, then the accuracy improves by a factor of 10, while if $\varepsilon_f = 10^{-12}$, then the accuracy improves by a factor of 100. Thus, we can expect improvements of one to two decimal places in the derivative at the cost of an additional function evaluation. The situation does not improve if we use a higher-order difference approximation since in this case ϕ is of the general form

$$\phi(h; \mu) = \mu_1 h^{2p} + \mu_2 \frac{\epsilon_f^2}{h^2},$$

where $p \ge 2$ is the order of the approximation and the constants μ_1 and μ_2 are determined by the approximation. Thus, the expected error is on the order of $\varepsilon_f^{p/(p+1)}$ but at the cost of p function evaluations.

Lemma 4.1 and Theorem 4.2 extend a result in [2] by establishing a lower bound for $\mathrm{E}\{\mathcal{E}_c\}$ and thus providing sharp estimates for the minimal value of $\mathrm{E}\{\mathcal{E}_c(h)\}$ as a function of h. These bounds show that the parameter h_M is nearly optimal in the sense that

$$\gamma_4\, \mu_L^{2/3}\, \varepsilon_f^{4/3} \le \mathrm{E}\{\mathcal{E}_c(h_M)\} \le \gamma_4\, \mu_M^{2/3}\, \varepsilon_f^{4/3}.$$

Thus h_M satisfies the same bounds as in Theorem 4.2. These bounds also imply that

$$\mathrm{E}\{\mathcal{E}_c(h_M)\} \le \left(\frac{\mu_M}{\mu_L}\right)^{2/3} \min_{|h|\le h_0} \mathcal{E}_c(h).$$

Hence, the error $\mathrm{E}\{\mathcal{E}_c(h_M)\}^{1/2}$ grows by a factor that depends on $(\mu_M/\mu_L)^p$, where $p = 1/3$. This is an improvement over Corollary 3.3, where $p = 1$.

We now consider the standard central difference approximation to f'' and show that in this case we again obtain expected errors of order $\varepsilon_f^{1/2}$. We express the results in terms of the error function

$$\mathcal{E}_2(h) = \left(\frac{f(t_0+h) - 2f(t_0) + f(t_0-h)}{h^2} - f_s''(t_0)\right)^2.$$

Lemma 4.3 *Assume that $f_s : \mathbb{R} \mapsto \mathbb{R}$ is four times differentiable in I, and set μ_L and μ_M to the minimum and maximum of $|f_s^{(4)}|$ on I, respectively. If $|h| \le h_0$, then*

$$\phi_2(h;\mu_L) \le \mathrm{E}\left\{\mathcal{E}_2(h)\right\} \le \phi_2(h;\mu_M),$$

where

$$\phi_2(h;\mu) = \left(\frac{\mu}{12}\right)^2 h^4 + 6\frac{\epsilon_f^2}{h^4}.$$

Proof. Taylor's theorem shows that

$$\frac{f_s(t_0+h) - 2f_s(t_0) + f_s(t_0-h)}{h^2} - f_s''(t_0) = \frac{1}{12} f_s^{(4)}(\xi) h^2$$

for some $\xi \in [t_0 - h, t_0 + h] \subseteq I$, and thus

$$\frac{f(t_0+h) - 2f(t_0) + f(t_0-h)}{h^2} - f_s''(t_0) = \frac{1}{12} f_s^{(4)}(\xi) h^2 + \frac{\varepsilon(t_0+h) - 2\varepsilon(t_0) + \varepsilon(t_0-h)}{h^2}.$$

The iid assumption on the noise $\varepsilon(\cdot)$ and (6) then yield that

$$\mathrm{E}\left\{\mathcal{E}_2(h)\right\} = \mathrm{E}\left\{\left(\frac{f(t_0+h) - 2f(t_0) + f(t_0-h)}{h^2} - f_s''(t_0)\right)^2\right\} = \left(\frac{1}{12} f_s^{(4)}(\xi)\right)^2 h^4 + 6\frac{\epsilon_f^2}{h^4}.$$

The result is a direct consequence of this expression. ■

We now determine the global minimizers of the mapping ϕ_2 in Lemma 4.3. A calculation shows that ϕ_2 is uniformly convex and that

$$\min_h \phi_2(h;\mu) = \gamma_6 \mu \varepsilon_f, \qquad \gamma_6 = \left(\tfrac{1}{6}\right)^{1/2} \approx 0.408. \tag{14}$$

Also note that the unconstrained global minimizers for μ_M are defined by

$$|h_M| = \gamma_7 \left(\frac{\varepsilon_f}{\mu_M}\right)^{1/4}, \qquad \gamma_7 = 2^{5/8}\, 3^{1/8} \approx 2.33. \tag{15}$$

The next result analyzes the behavior of $\mathrm{E}\left\{\mathcal{E}_2(h)\right\}$ and shows that h_M plays an important role in this analysis. This result follows the pattern used for the forward difference approximation to f' in Section 3.

Theorem 4.4 *Assume that $f_s : \mathbb{R} \mapsto \mathbb{R}$ is four times differentiable in I; set μ_L and μ_M to the minimum and maximum of $|f_s^{(4)}|$ on I, respectively; and define μ_M by (15). If $|h_M| \le h_0$, then*

$$\gamma_6 \mu_L \varepsilon_f \le \min_{|h| \le h_0} \mathrm{E}\left\{\mathcal{E}_2(h)\right\} \le \gamma_6 \mu_M \varepsilon_f.$$

Proof. Lemma 4.3 implies that

$$\min_{|h| \le h_0} \phi_2(h;\mu_L) \le \min_{|h| \le h_0} \mathrm{E}\left\{\mathcal{E}_2(h)\right\} \le \min_{|h| \le h_0} \phi_2(h;\mu_M).$$

The proof now follows from the observations that $\gamma_6\, \mu\varepsilon_f$ is a lower bound for $\phi_2(\cdot, \mu)$ and that (14) holds if $\mu = \mu_M$ and $|h_M| \le h_0$. ■

The similarity between Theorem 4.4 for f_s'' and Theorem 3.2 for f_s' is of interest. In both cases the bounds depend linearly on $\mu_L\varepsilon_f$ and $\mu_M\varepsilon_f$. In this vein, note that Theorem 4.4 and (14) imply that if h_M is defined by (15), then

$$\gamma_6\mu_L\varepsilon_f \le \mathrm{E}\left\{\mathcal{E}_2(h_M)\right\} \le \gamma_6\mu_M\varepsilon_f, \qquad \mathrm{E}\left\{\mathcal{E}_2(h_M)\right\} \le \left(\frac{\mu_M}{\mu_L}\right) \min_{|h|\le h_0} \mathrm{E}\left\{\mathcal{E}_2(h)\right\}.$$

These inequalities show that h_M yields nearly optimal bounds.

Similar results have appeared in the literature. If the computed function f is deterministic and of the form (9) where the error function satisfies (10), then [5, 6] show that the parameter h that minimizes a bound on the l_1 error for the centered approximations to f' and f'' are, respectively,

$$\left(\frac{3\varepsilon_A}{|f^{(3)}(\xi_a)|}\right)^{1/3} \quad \text{and} \quad 2\left(\frac{3\varepsilon_A}{|f^{(4)}(\xi_b)|}\right)^{1/4}$$

for some ξ_a and ξ_b. As discussed in Section 3, these results do not guarantee near optimality and rely on the bound ε_A instead of the noise level ε_f.

5 An Algorithm for Forward Difference Estimates

In this section we summarize our algorithm for obtaining a nearly optimal estimate of the derivative for forward differences. A similar approach can be applied for other difference schemes. We also indicate how far from optimality our estimates could be in practice.

As in the previous sections, $f : \mathbb{R} \mapsto \mathbb{R}$ is the computed function and is defined in an interval I around t_0. However, since our main interest is in deterministic simulations, our algorithms refer only to the computed function f. Estimates of f_s'' required by the theoretical results are replaced by estimates of f''.

We assume that a positive estimate of the noise level ε_f is available. We obtain ε_f with the ECnoise algorithm detailed in [13], at an expense of 6–8 function evaluations. In our experience, these function values can be saved by reusing noise estimates obtained at a base point $t_1 \ne t_0$. This is almost certainly the case if t_0 is of the same order of magnitude as t_1. If the relative noise is expected to be constant, then the scaled noise $(f(t_0)/f(t_1))\, \varepsilon_f$ can be used.

We next require a coarse estimate of the second derivative. The step h_M in (8) requires a bound, $\mu_M = \max_{t\in I} |f''(t)|$; but in practice this is unavailable, and we rely on an estimate μ of $|f''(t_0)|$. Given the estimate μ and the noise level ε_f, we use

$$h^* = 8^{1/4}\left(\frac{\varepsilon_f}{\mu}\right)^{1/2} \qquad (16)$$

as the difference parameter to estimate $f'(t_0)$. Below we show that h^* is relatively insensitive to the estimate $\mu \approx |f''(t_0)|$, and thus a rough estimate of f'' yields a highly accurate approximation of f'. We also note that it may be possible to estimate $|f''(t_0)|$ from the function values used by ECnoise, but we do not pursue this possibility here.

Our aim is to produce an estimate of f'' with a few function evaluations. Our algorithm for obtaining the estimate relies on the second-order difference

$$\Delta(h) = |f(t_0 - h) - 2f(t_0) + f(t_0 + h)|.$$

We do not want to choose h too small since then computational noise corrupts the computation of $\Delta(h)$, while if h is too large, then truncation errors destroy the accuracy of $\Delta(h)$. We accept h as being sufficiently large if

$$\frac{\Delta(h)}{\varepsilon_f} \geq \tau_1 \tag{17a}$$

for some constant $\tau_1 \gg 1$. This signal-to-noise test relies on the observation that if h is too small, then $\Delta(h)$ is of the order of ε_f. We also require that all function values that contribute to $\Delta(h)$ be similar and thus that h not be too large. We thus require that

$$|f(t_0 \pm h) - f(t_0)| \leq \tau_2 \max\{|f(t_0)|, |f(t_0 \pm h)|\} \tag{17b}$$

for some constant $\tau_2 \in (0, 1)$. This test guarantees that there is cancellation in the computation of $\Delta(h)$ and thus that h is not too large.

If h satisfies the tests in (17), then $\mu = \Delta(h)/h^2$ should be a suitable estimate of $|f''(t_0)|$. We could use a search technique to determine an h that satisfies these conditions, but we prefer to use at most 2–4 function evaluations to this task. Our heuristics for choosing h are outlined in Figure 5.1. These heuristics are closely related to the work [5] on determining suitable difference parameters for estimating $|f''|$ by accepting $\mu = \Delta(h)/h^2$ as an estimate if

$$\frac{\Delta(h)}{4\varepsilon_A} \in \left[10, 10^3\right].$$

Our computational results show that the lower bound in this test should be replaced by (17a) with $\tau_1 = 100$, while the upper bound is too restrictive. We often accept h values that violate this upper bound but satisfy (17b) with $\tau_2 = 0.1$.

Algorithm 5.1 ***Heuristic for an estimate of*** $|f''(t_0)|$**.**

> *Set* $\tau_1 = 100$ *and* $\tau_2 = 0.1$ *in* (17).
>
> *Set* $h_a = \varepsilon_f^{1/4}$ *and* $\mu_a = \Delta(h_a)/h_a^2$. *Exit with* $\mu = \mu_a$ *if* h_a *satisfies* (17).
>
> *Set* $h_b = (\varepsilon_f/\mu_a)^{1/4}$ *and* $\mu_b = \Delta(h_b)/h_b^2$. *Exit with* $\mu = \mu_b$ *if* h_b *satisfies* (17).
>
> *Exit with* $\mu = \mu_b$ *if* $|\mu_a - \mu_b| \leq \frac{1}{2}\mu_b$.

The choice of $h_a = \varepsilon_f^{1/4}$ in Algorithm 5.1 is motivated by (15). Theorem 4.4 shows that this choice of h_a is likely to be acceptable if μ_M in (15) is of modest size. If μ_M is large, then the estimate of μ_a tends to be large, and thus we use this estimate in our choice of h_b. In our computational experiments $h_b \ll h_a$, but this is not guaranteed. In most cases the choice of h_b produces a good estimate of $|f''(t_0)|$, but this estimate may not satisfy condition (17) when the noise is relatively large. Thus Algorithm 5.1 also sets $\mu = \mu_b$ if both estimates μ_a and μ_b are similar.

It may be the case, as will be seen in the numerical experiments, that after running this algorithm we are unable to satisfy both tests, though this situation generally occurs when the noise is relatively large. One should proceed carefully in such an event. The subject of characterizing when the function is too noisy to allow for estimates of f'' or f' is beyond the scope of this paper.

Having defined h^* in (16) (when an estimate μ is computable) as an approximation to the h_M desired in theory, we now study the sensitivity of this step to ε_f and μ_M.

Theorem 5.2 *Assume that $f_s : \mathbb{R} \mapsto \mathbb{R}$ is twice differentiable in I; set μ_L and μ_M to the minimum and maximum of $|f_s''|$ on I, respectively; and define h_M by (8). If $\alpha h_M \le h_0$, then*

$$\frac{\mathrm{E}\left\{\mathcal{E}(\alpha h_M)\right\}}{\mathrm{E}\left\{\mathcal{E}(h_M)\right\}} \le \frac{\mu_M^2}{\mu_L^2 + \mu_M^2}\left(\frac{1}{\alpha^2} + \alpha^2\right).$$

Proof. Lemma 3.1 shows that for any $h_\alpha \le h_0$ there is an $\mu_\alpha \le \mu_M$ such that

$$\mathrm{E}\left\{\mathcal{E}(h_\alpha)\right\} = \frac{1}{4}\mu_\alpha^2 h_\alpha^2 + 2\frac{\varepsilon_f^2}{h_\alpha^2}.$$

Setting $h_\alpha = \alpha h_M$ and using the definition of h_M in this expression, we have that

$$\mathrm{E}\left\{\mathcal{E}(\alpha h_M)\right\} = \frac{\varepsilon_f}{\sqrt{2}}\left(\frac{\mu_\alpha^2\alpha^2}{\mu_M} + \frac{\mu_M}{\alpha^2}\right) = \frac{\varepsilon_f}{\sqrt{2}}\left(\frac{\mu_\alpha^2\alpha^4 + \mu_M^2}{\mu_M\alpha^2}\right).$$

We now finish the proof by noting that

$$\frac{\mathrm{E}\left\{\mathcal{E}(\alpha h_M)\right\}}{\mathrm{E}\left\{\mathcal{E}(h_M)\right\}} = \frac{1}{\alpha^2}\left(\frac{\mu_\alpha^2\alpha^4 + \mu_M^2}{\mu_1^2 + \mu_M^2}\right) \le \frac{\mu_M^2}{\mu_L^2 + \mu_M^2}\left(\frac{1}{\alpha^2} + \alpha^2\right),$$

where the upper bound is obtained by using μ_L as a lower bound for μ_1 and μ_M as an upper bound for μ_α. ■

Theorem 5.2 corrects the bound in the first part of Theorem 5.1 in [2], which depends on the claim that (in our notation)

$$\mathrm{E}\left\{\mathcal{E}(h_M)\right\} = \sqrt{2}\mu_M\varepsilon_f.$$

Simple examples show that this claim and the result in [2] does not hold for nonquadratic functions. In this work we use the weaker result

$$\frac{\mathrm{E}\left\{\mathcal{E}(\alpha h_M)\right\}}{\mathrm{E}\left\{\mathcal{E}(h_M)\right\}} \le \left(\frac{1}{\alpha^2} + \alpha^2\right),$$

since this bound quantifies the effects of estimating the unknown quantities ε_f and μ_M in h^*. If $h^* = \alpha h_M$, then

$$\mathrm{E}\left\{\mathcal{E}(h^*)\right\}^{1/2} \le \left(\frac{1}{\alpha^2} + \alpha^2\right)^{1/2} \mathrm{E}\left\{\mathcal{E}(h_M)\right\}^{1/2} \le \left(\frac{1}{\alpha} + \alpha\right) \mathrm{E}\left\{\mathcal{E}(h_M)\right\}^{1/2},$$

and thus the expected error $\mathrm{E}\left\{\mathcal{E}(h^*)\right\}^{1/2}$ grows slowly with the ratio $\alpha = h^*/h_M$. This result shows, in particular, that if we misestimate ε_f/μ_M by a factor of 10^d for some $d > 0$, then $\alpha \in [10^{-d/2}, 10^{d/2}]$, and thus

$$\mathrm{E}\left\{\mathcal{E}(h^*)\right\}^{1/2} \le 10^{d/2} \mathrm{E}\left\{\mathcal{E}(h_M)\right\}^{1/2}.$$

Thus, we are likely to lose $d/2$ digits relative to the best possible value. We will present computational verification of this result in the next section.

6 Computational Experiments

In this section we assume that $f : \mathbb{R}^n \mapsto \mathbb{R}$ and analyze the empirical properties and limitations of the algorithm proposed in Section 5 on a variety of problems. In all cases we choose a direction $p \in \mathbb{R}^n$ and apply Algorithm 5.1 to the function

$$\phi(t) = f(x_0 + tp)$$

for some $x_0 \in \mathbb{R}^n$. We generate a difference parameter h^* and compare the resulting difference estimate with estimates obtained from other h values against an accurate directional derivative $f'(x_0; p)$ by examining the errors

$$\mathcal{E}(h) = \left(\frac{f(x_0 + hp) - f(x_0)}{h} - f'(x_0; p)\right)^2.$$

This approach depends on choosing a value for $f'(x_0; p)$ that is itself noisy. In the case of the stochastic functions in Figure 3, we used the analytic derivative of the mean, but this approach will not work for deterministic functions.

When reliable hand-coded derivatives are not available, the approach we take is to set $f'(x_0; p)$ to the directional derivative obtained through automatic differentiation. We have found stable estimates of the derivatives for the problems presented here using the MATLAB tool IntLab [14]. In our experience IntLab produces estimates of the derivatives that are within the noise level of the derivative and are thus as accurate as possible.

We have already seen an example of our approach for the Higham function in Figure 2 (right) where the relative error

$$\mathcal{R}(h) = \frac{1}{|f'(x_0; p)|} \sqrt{\mathcal{E}(h)} = \frac{1}{|f'(x_0; p)|} \left| \frac{f(x_0 + hp) - f(x_0)}{h} - f'(x_0; p) \right|$$

was plotted for many different h. Despite the fact that the noise in this function is by no means iid, the h^* estimated by our algorithm falls very close to the minimum of the errors.

Stochastic Problems

We have seen that our computed h^* is close to the optimal h for the stochastic polynomials in Figure 3. As an example of a more complex problem where the noise level is not specified in advance, we consider the `MCfinance` problem

$$f(x) = (2\pi)^{-n/2} \int_{\mathbb{R}^n} \prod_{i=0}^{n} \frac{1}{1 + r_i(u,x)} e^{-\frac{\|u\|^2}{2}} du$$

discussed in [13]. We set $n = 3$ and generate 10^6 standard normal random variables to obtain a Monte Carlo evaluation of f at $x_0 = [.1, .1, .1]$. For a random direction p, the directional derivative $f'(x_0; p)$ is computed with IntLab.

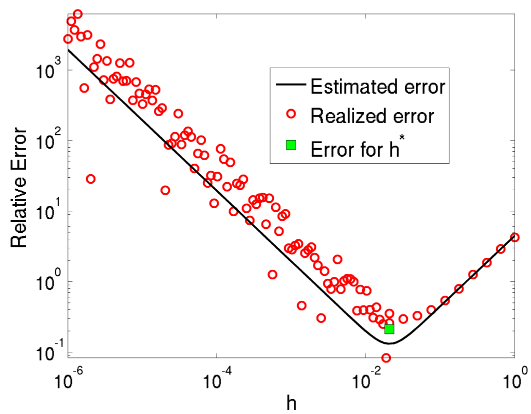

Figure 4: Plot of realizations of $\mathcal{R}(h)$ for the stochastic MCfinance problem.

While the majority of the estimates of $f'(x_0; p)$ had relative errors $\mathcal{R}(h) > 1$, Figure 4 shows that our h^* is able to obtain the correct first digit of f'. This shows that proper choice of the difference parameter h can yield (admittedly coarse) estimates of f' for problems for which derivative estimation is difficult at best.

We are also interested in determining the difference between the realized error $\mathcal{R}(h)$ for a given h and the expected error $\mathrm{E}\{\mathcal{E}(h)\}$ as defined in Section 3. Since Lemma 3.1 shows that the expected error can be estimated by

$$\mathcal{R}_e(h) = \frac{1}{|f'(x_0; p)|} \sqrt{\frac{1}{4}\mu^2 h^2 + 2\frac{\varepsilon_f^2}{h^2}}, \tag{18}$$

where μ is our estimate of $f''(x_0; p)$, we have plotted this estimate in Figure 4. As can be seen in this figure, $\mathcal{R}_e(h)$ is close to the actual errors. Also, the value of h^* produced by our algorithm is close to the minimum of the estimates $\mathcal{R}_e(h)$ of the errors.

Noisy Quadratic Problems

As an initial example of deterministic functions, we consider the noisy quadratic function defined by $f(b) = \|x(b)\|^2$ where $x(b)$ is the bicgstab solution of $Ax = b$ using a relative tolerance of $\tau = 10^{-3}$.

We use the same set of 116 matrices used in [13], representing all symmetric positive definite matrices of dimension less than 10^4 in the University of Florida (UF) Sparse Matrix Collection [4]. Following [13], we scale the matrices by their diagonals and randomly select the base point x_0 and direction p. The resulting matrices have an (absolute) noise level ranging in order from 10^{-16} to 10^0. The directional derivative $f'(x_0; p)$ is computed with IntLab.

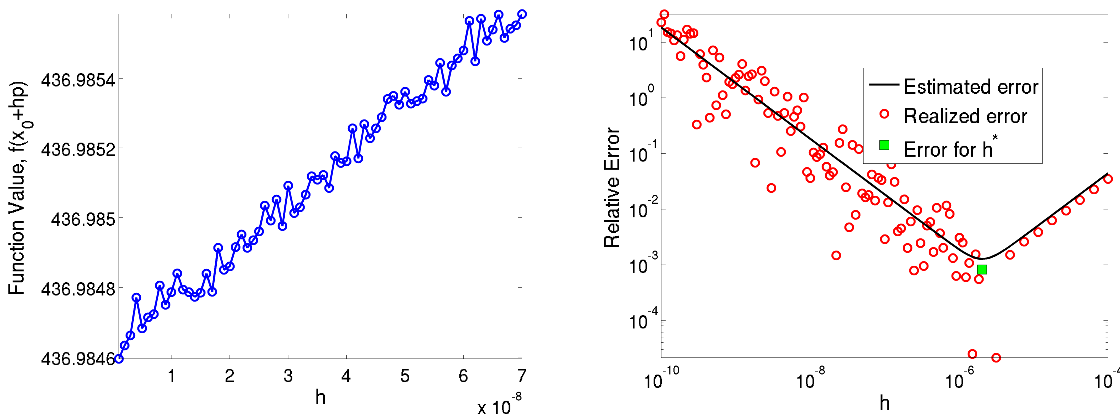

Figure 5: Quadratic associated with the UF matrix bcsstk26 ($n = 1922$): Plot of the mapping $h \mapsto f(x_0 + hp)$ (left) and plot of the relative errors $\mathcal{R}$ (right).

Figure 5 (left) illustrates the noise in the function values for the UF matrix bcsstk26. The noise produced for this matrix is typical for the bicgstab algorithm. Figure 5 (right) shows realizations of the relative error $\mathcal{R}(h)$ for many different h values as well as the expected error as estimated by (18) for the quadratic function associated with the bcsstk26 matrix. Plots of the relative error $\mathcal{R}$ for the other matrices in the UF collection are similar. The behavior of $\mathcal{R}$ in this plot bears a striking resemblance to that for the stochastic quadratic in Figure 3 (left), illustrating that the proposed algorithm can also work well for deterministic noise.

We ran our algorithm on the entire set of 116 matrices and found that we were unable to obtain reliable estimates of $|f''(x_0)|$ using the heuristic in Algorithm 5.1 on 16 of the matrices. These 16 cases generally correspond to the noisiest problems, for which we do not expect our algorithm to produce stable estimates.

Figure 6 summarizes the results for the remaining 100 quadratics, when the problems are sorted by the relative errors $\mathcal{R}(h^*)$. For comparison, we also show the realized error for difference parameters that are two orders of magnitude smaller and larger than h^*. As predicted by Theorem 5.2, the general trend is that h^* yields derivatives with two more correct digits than $10^{\pm 2}h^*$. For five of the 100 matrices we see that smaller errors are realized for one of the alternatives, but these improvements tend to be minor and reflect the variability expected from Figure 5.

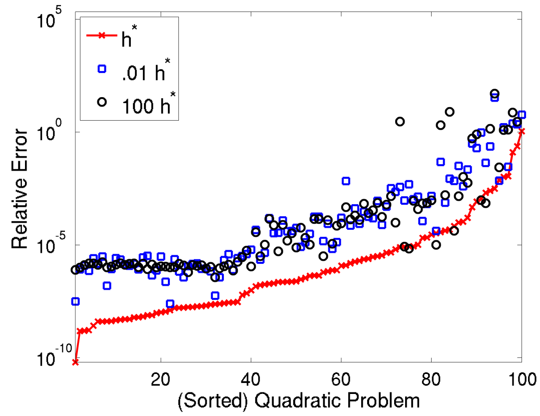

Figure 6: Plot of $\mathcal{R}(h^*)$ and $\mathcal{R}(10^{\pm 2}h^*)$ for 100 quadratics.

Smooth Problems

A feature of quadratics, provided they are not too noisy, is that the second derivative remains relatively constant, thus allowing reasonable estimates of the second derivative to be obtained for any h that is not small. Figure 5 (left) shows that for these problems we need to choose $h \gg 10^{-8}$. We now examine how the algorithm performs on highly nonlinear problems, where this is not the case.

We analyze the same five MINPACK-2 problems considered in [11] but use more dimensions. In [11] $n = 10,000$, while we use the five dimensions

$$n \in \{2.5,\ 10,\ 40,\ 160,\ 640\} \cdot 10^3$$

for each problem. In addition to representing high-dimensional nonlinear functions, these problems have the benefit of possessing hand-coded derivatives, which we use to obtain $f'(x_0; p)$ for our comparisons. In support of the reliability of IntLab for the other test problems, Table 1 shows that the IntLab derivatives agree remarkably well with the hand-coded directional derivatives for the MINPACK-2 problems. Table 1 also shows that, as expected, the hand-coded derivatives are faster and that the computing time for both types of derivatives grows linearly with n.

Table 1: Comparison of IntLab and hand-coded derivatives on the ept problem.

n	Relative Error	Hand-coded time [s]	IntLab time [s]
2500	$3.17 \cdot 10^{-15}$	$2.74 \cdot 10^{-4}$	5.36
10000	$5.53 \cdot 10^{-15}$	$9.90 \cdot 10^{-4}$	22.2
40000	$2.79 \cdot 10^{-14}$	$4.23 \cdot 10^{-3}$	109

Figure 7 summarizes the results for these 25 problems when the problems are sorted by the relative errors $\mathcal{R}(h^*)$. As before, we plot the realized error for difference parameters that are two orders of magnitude perturbations of h^* and note that the error from h^* is smaller than these two alternatives for all 25 problems. These results show that our techniques are applicable to problems with a large number of variables and that the results are independent of the number of variables. For the specific case of the ept problem with $n = 640,000$, Figure 7 (right) shows that the relative error $\mathcal{R}(h)$ generally behaves as before but that the realizations no longer agree well with the estimated error for small h. In this case, the noise level ε_f has been overestimated roughly by a factor 6; but as predicted by Theorem 5.2, the resulting h^* is insensitive to this small misestimation.

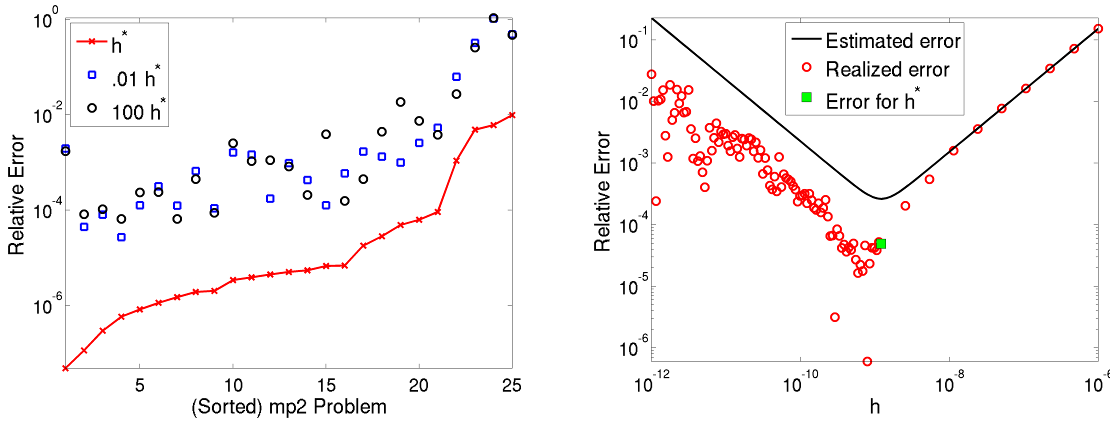

Figure 7: Plot of $\mathcal{R}(h^*)$ and $\mathcal{R}(10^{\pm 2}h^*)$ for 25 MINPACK-2 problems (left) and $\mathcal{R}(h)$ for problem ept with $n = 640,000$ (right).

References

[1] C. H. Bischof, H. M. Bücker, P. Hovland, U. Naumann, and J. Utke, eds., *Advances in Automatic Differentiation*, vol. 64 of Lecture Notes in Computational Science and Engineering, Springer, 2008.

[2] R. C. M. Brekelmans, L. Driessen, H. J. M. Hamers, and D. den Hertog, *Gradient estimation schemes for noisy functions*, Journal of Optimization Theory and Applications, 126 (2005), pp. 529–551.

[3] H. M. Bücker, G. Corliss, P. Hovland, U. Naumann, and B. Norris, eds., *Automatic Differentiation: Applications, Theory, and Implementations*, vol. 50 of Lecture Notes in Computational Science and Engineering, Springer, 2006.

[4] T. Davis, *The University of Florida sparse matrix collection*, manuscript, University of Florida, 2008. Submitted to ACM Transactions on Mathematical Software.

[5] P. E. Gill, W. Murray, M. A. Saunders, and M. H. Wright, *Computing forward-difference intervals for numerical optimization*, SIAM J. Sci. Statist. Comput., 4 (1983), pp. 310–321.

[6] P. E. Gill, W. Murray, and M. H. Wright, *Practical Optimization*, Academic Press, 1981.

[7] A. Griewank and A. Walther, *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*, SIAM, 2008.

[8] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, SIAM, 2002.

[9] A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker, and C. S. Woodward, *SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers*, ACM Transactions on Mathematical Software, 31 (2005), pp. 363–396.

[10] D. Knoll and D. Keyes, *Jacobian-free Newton–Krylov methods: A survey of approaches and applications*, Journal of Computational Physics, 193 (2004), pp. 357–397.

[11] C.-J. Lin and J. J. Moré, *Newton's method for large bound-constrained optimization problems*, SIAM Journal on Optimization, 9 (1999), pp. 1100–1127.

[12] J. N. Lyness, *Has numerical differentiation a future?*, in Proceedings Seventh Manitoba Conference on Numerical Mathematics, Utilitas Mathematica Publishing, 1977.

[13] J. J. Moré and S. M. Wild, *Estimating computational noise*, Preprint ANL/MCS-P1721-0210, Argonne National Laboratory, Argonne, Illinois, 2010.

[14] S. M. Rump, *INTLAB - INTerval LABoratory*, in Developments in Reliable Computing, T. Csendes, ed., Kluwer Academic Publishers, Dordrecht, 1999, pp. 77–104.