ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, Illinois 60439

Do You Trust Derivatives or Differences?*

Jorge J. Moré and Stefan M. Wild

Mathematics and Computer Science Division

Preprint ANL/MCS-P2067-0312

April 2012

*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.

Do You Trust Derivatives or Differences?*

Jorge J. Moré and Stefan M. Wild

Abstract

We analyze the relationship between the noise level of a function and the accuracy and reliability of derivatives and difference estimates. We derive and empirically validate measures of quality for both derivatives and difference estimates. Using these measures, we quantify the accuracy of derivatives and differences in terms of the noise level of the function. An interesting consequence of these results is that the derivative of a function is not likely to have working precision accuracy for functions with modest levels of noise.

1 Introduction

Although the accurate estimation of the derivative of a nonlinear mapping $f : \mathbb{R}^n \mapsto \mathbb{R}$ is a central problem in computational science, most of the literature in this area is restricted to the case where the noise level in the function f is on the order of machine precision. In this work we study the relationship between the noise of f and the accuracy of (finite-precision) estimates f' of the derivative and estimates δf obtained by differences of functions values.

In [13] we defined computational noise for functions determined by a stochastic process. We showed that the stochastic framework can be applied to deterministic functions with results that are consistent with our intuition. In addition, we have developed the ECNoise algorithm that determines the noise level ε_f of the function near a base point with a few additional function evaluations. Consult [13] for an extensive discussion of noise and computational results for the ECNoise algorithm.

The noise level ε_f of a function f measures the accuracy and uncertainty in the computed function values. Thus it is important to know how the noise level of the function affects the accuracy of f' and δf. We address two basic questions:

- ◇ Can we rely on f' and δf when the noise level ε_f of f is high?
- ◇ What is the relationship between ε_f and the accuracy of f' and δf?

As we will see, answering these questions requires the development of tools for measuring the accuracy of f' and δf and yields surprising results.

A systematic approach for producing accurate estimates of the derivative of a function f is to use automatic differentiation tools [1, 8, 14]. The derivative codes produced by these tools can be shown to be backward stable [8, Section 3.4] under suitable conditions on the elemental functions in the function evaluation, and thus these tools are frequently claimed to produce derivatives that are accurate to working precision. The claim [8, page 50] that

*Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439.

"once an AD tool has been applied successfully ... we have never heard of complaints concerning the precision of the derivative values"

is typical. We will lend empirical evidence supporting this claim when the noise level of the function is on the order of machine precision, but we also show that this claim fails for functions with modest levels of noise.

The derivative produced by the forward mode of automatic differentiation tools is essentially a line-by-line differentiation of the function evaluation code and is thus representative of the default derivative code that could be produced by a computational scientist. On the other hand, a computational scientist can use additional information to produce a derivative with different numerical characteristics; we emphasize that our conclusions about the performance of derivative codes may not apply to these alternative hand-coded derivatives.

The accuracy of a difference estimate δf of the derivative is well known to depend on the choice of difference parameter. We have shown [12] that we can determine a difference parameter that is provably near-optimal when the function values are determined by a stochastic process. We have also shown that the stochastic techniques can be used to determine accurate derivative estimates when the function is deterministic. Additionally, we have shown that for a forward difference estimate the accuracy of δf is closely related to $\varepsilon_f^{1/2}$.

Our computational results use a basic component in scientific simulations: the iterative solution of systems of linear equations with a Krylov solver. Section 2 examines the continuity and differentiability properties of mappings defined by a Krylov solver. We show that given a continuously differentiable function $h : \mathbb{R}^n \mapsto \mathbb{R}^n$, if there is an upper bound m on the number of iterations and if $y_\tau : \mathbb{R}^n \mapsto \mathbb{R}^n$ is defined by letting $y_\tau(x)$ be the approximate solution of $Ay = h(x)$ computed by a Krylov solver with the standard relative residual termination test with tolerance τ, then y_τ is continuously differentiable in any compact set Ω of starting points except for a finite number of tolerances τ. This result guarantees the existence of the derivative for the mappings considered in this work. We also show that y'_τ is uniformly bounded in Ω for all $\tau \geq \tau_{\min}$, where $\tau_{\min} \geq 0$ depends on the bound m on the number of iterations and the properties of the iterative method. For example, if the Krylov solver is guaranteed to terminate within m iterations, then $\tau_{\min} = 0$.

Results on the differentiability properties of mappings defined by iterative methods have appeared in the literature [3, 4, 7], but these results assume that the mapping of interest is the limit of a sequence of differentiable mappings. In particular, no termination test is performed. In this work we are interested mainly in the differentiability properties of y_τ for finite values of $\tau > 0$ and are only marginally concerned with the limit of y_τ as $\tau \to 0$.

Section 3 examines the relationship between the tolerance τ and the relative noise level of functions f_τ that depend on τ. Our computational experiments use functions defined by $f_\tau(x) = \|y_\tau(x)\|^2$, where $y_\tau(x)$ has already been defined as the approximate solution of $Ay = h(x)$ obtained by a Krylov solver and where the matrices A are the symmetric positive definite matrices of dimension less than 10^4 in the University of Florida Sparse Matrix Collection [6]. We discuss results for a variety of solvers and tolerances but focus on bicgstab; results for other solvers, including pcg, minres, and gmres, appear in [11]. The

main observation in this section is that the dependence of the noise level on the tolerance follows a general pattern: The noise level is low for large tolerances but increases sharply as the tolerance τ decreases and then decreases gradually as τ decreases further. The relative noise level of at least 80% of the functions is higher than 10^{-10} for at least one tolerance. This observation is important because it shows that we can obtain at most 10 significant digits from these functions.

We study the accuracy of the derivative of f_τ in Section 4 in terms of the relative error metric (see, e.g., Ziv [17]) between the derivatives computed with the IntLab [14] and ADiMat [1] `Matlab`-based automatic differentiation tools and the relative error between two nearby derivative values. We show that the relative error of f'_τ is usually on the order of the noise level ε_f. This observation shows that we cannot expect to obtain double precision accuracy if the functions have modest levels of noise. A surprising outcome of these experiments is that the derivatives of functions with a high noise level are unstable.

Section 5 reviews techniques [12] for determining a nearly optimal difference parameter h^* for a forward difference estimate δf_τ and then presents computational results concerning the accuracy of δf_τ as measured by the relative error metric $\mathsf{re}(\delta f_\tau, f'_\tau)$ between δf_τ and f'_τ. The main contribution of this section is a computable estimate $\Gamma(\delta f_\tau)$ of the relative error $\mathsf{re}(\delta f_\tau, f'_\tau)$ determined from the same information needed to determine h^*. Our numerical results show remarkable agreement between $\Gamma(\delta f_\tau)$ and $\mathsf{re}(\delta f_\tau, f'_\tau)$, and thus we are able to predict the accuracy of δf_τ as a by-product of the computation of h^*.

The results in Sections 4 and 5 show that the derivative estimate f'_τ obtained from the automatic differentiation tools have higher accuracy than δf_τ when the function has modest levels of noise. However, these results also show that if the noise level is high, then this conclusion can fail to hold. We study these functions in Section 6 with the tools that we have developed to estimate the accuracy of derivative and difference estimates; the aim is to analyze cases where the accuracy of f'_τ and δf_τ is not acceptable. The results for the bicgstab and idr solvers are of particular interest because these solvers generate problems with unusually high levels of noise.

2 Continuity and Differentiability of Krylov Solvers

We begin by investigating the continuity properties of functions defined by the approximate solution of a system of linear equations by an iterative process. We consider Krylov solvers because these solvers are commonly used in computational science applications, but we expect to find similar behavior for other solvers and other problems.

Let $h : \mathbb{R}^n \mapsto \mathbb{R}^n$ be a continuous mapping, and let $y_\tau : \mathbb{R}^n \mapsto \mathbb{R}^n$ be the mapping defined as the solution of $Ay = h(x)$ returned by a deterministic Krylov solver that uses a tolerance $\tau > 0$. We use the standard termination test based on the residual,

$$\|Ay - h(x)\| \leq \tau \|h(x)\|, \tag{2.1}$$

where $\|\cdot\|$ is the l_2 norm. The matrix A is assumed to be nonsingular.

The iterates of a Krylov solver for the system $Ay = h(x)$ depend on x through the starting point. We assume that the Krylov solver has no breakdowns and that it generates iterates $y_0(x), \ldots, y_m(x)$ until the termination test (2.1) is satisfied or the Krylov solver reaches a limit m on the number of iterations. We assume that the starting point $y_0(x)$ is a continuously differentiable function of x but that m is independent of x. We could choose m to be a fixed number or to be dependent on the number of variables, for example, $m = 2n$. Also note that we can always assume that the Krylov solver generates m iterates by defining $y_k(x) = y_j(x)$ for $j < k \le m$ if $y_j(x)$ satisfies (2.1) for some $j < m$.

The function $y_k : \mathbb{R}^n \mapsto \mathbb{R}^n$ maps x into the kth iterate generated by the solver. Each function y_k is continuously differentiable because the iterates are the composition of continuously differentiable functions of the starting point $y_0(x)$. A glance at the formulation of any Krylov solver confirms this claim, but this claim also holds for most iterative solvers. See, for example, the implementations detailed in [2, 15, 16].

The mapping $y_\tau : \mathbb{R}^n \mapsto \mathbb{R}^n$ is defined by letting $y_\tau(x)$ be the first iterate that satisfies the termination test (2.1). The properties of y_τ in a neighborhood of a point x_0 with $h(x_0) \neq 0$ can be described in terms of the parameters

$$\tau_k = \frac{\|Ay_k(x_0) - h(x_0)\|}{\|h(x_0)\|}, \qquad 0 \le k \le m. \tag{2.2}$$

The definition of τ_k shows that the kth iterate of the solver satisfies the termination test (2.1) for $x = x_0$ if and only if $\tau \ge \tau_k$. Hence, the mapping y_τ is well defined at x_0 if $\tau \ge \tau_{\min}$, where

$$\tau_{\min} = \min \{\tau_k : 0 \le k \le m\}.$$

We assume that $\tau \ge \tau_{\min}$ in our analysis of the properties of y_τ because (2.1) cannot be satisfied within m iterations when $\tau < \tau_{\min}$. Later we will assume that the Krylov solver terminates at the solution in at most m iterations and thus $\tau_{\min} = 0$, but this situation may not happen if m is not large enough or if the properties of the matrix A do not guarantee termination of the solver. The following result addresses the differentiability properties of y_τ in a neighborhood of x_0.

Theorem 2.1. *Assume that $h : \mathbb{R}^n \mapsto \mathbb{R}^n$ is continuously differentiable in a neighborhood of x_0 and that $h(x_0) \neq 0$. If $\tau \ge \tau_{\min}$, then the mapping $y_\tau : \mathbb{R}^n \mapsto \mathbb{R}^n$ is continuously differentiable in a neighborhood $N_\tau(x_0)$ of x_0 if $\tau \notin \{\tau_0, \ldots, \tau_m\}$, where τ_k is defined by (2.2). Moreover,*

$$y'_\tau(x) \in \left\{y'_k(x) : 0 \le k \le m\right\}, \qquad x \in N_\tau(x_0). \tag{2.3}$$

Proof. We have already noted that if $\tau \ge \tau_{\min}$, then y_τ is well defined at x_0. If the solver terminates at the kth iteration, then

$$\|Ay_k(x) - h(x)\| < \tau\|h(x)\|, \qquad \|Ay_j(x) - h(x)\| > \tau\|h(x)\|, \quad 0 \le j < k \tag{2.4}$$

holds for $x = x_0$ because $\tau \neq \tau_k$ and $h(x_0) \neq 0$. The continuity of h and y_j then show that (2.4) holds for all x in a neighborhood $N_\tau(x_0)$ of x_0. Hence, $y_\tau(x) = y_k(x)$ on $N_\tau(x_0)$, and this implies that y_τ is continuously differentiable in $N_\tau(x_0)$. In particular, $y'_\tau(x) = y'_k(x)$ on $N_\tau(x_0)$, and thus (2.3) holds. □

Another way to phrase the main conclusion of Theorem 2.1 is that y'_τ is continuous in a neighborhood of any point x_0 for almost all $\tau \geq \tau_{\min}$, and thus computing approximations to $y'_\tau(x_0)$ is well defined. In general, the terminology *for almost all* x means that the property in question holds except for a finite number of values of the parameter x.

Theorem 2.1 is local but can be extended by compactness arguments. Assume that a compact set $\Omega \subset \mathbb{R}^n$ contains all points of interest. Since y_τ is continuously differentiable in a neighborhood of any $x_0 \in \Omega$ except for τ in $\{\tau_0, \ldots, \tau_m\}$, we can cover Ω with a finite number of such neighborhood. Thus y_τ is continuously differentiable in Ω for

$$\tau \geq \tau_{\min} \equiv \max_{x \in \Omega} \left\{ \min_{0 \leq k \leq m} \left\{ \frac{\|Ay_k(x) - h(x)\|}{\|h(x)\|} \right\} \right\} \tag{2.5}$$

except for a finite number of τ. We formally state this result for later reference.

Theorem 2.2. *Assume that $h : \mathbb{R}^n \mapsto \mathbb{R}^n$ is continuously differentiable in a compact set Ω and that $h(x) \neq 0$ in Ω. If $\tau \geq \tau_{\min}$, where $\tau_{\min}$ is defined by (2.5), then the mapping $y_\tau : \mathbb{R}^n \mapsto \mathbb{R}^n$ is continuously differentiable in Ω for almost all values of τ. Moreover, y'_τ is uniformly bounded in Ω for almost all $\tau \geq \tau_{\min}$.*

Proof. We have already shown that y_τ is continuously differentiable for almost all τ. The proof of the uniform boundedness of y'_τ follows from (2.3) because this inclusion implies that

$$\|y'_\tau(x)\| \leq \max_{0 \leq k \leq m} \left\{ \|y'_k(x)\| \right\}, \qquad x \in \Omega,$$

whenever $y'_\tau(x)$ exists. The term on the right side is bounded, independent of τ, and thus y'_τ is uniformly bounded for almost all τ. □

In these results we have shown that y_τ is defined for $x \in \Omega$ with $\tau \geq \tau_{\min}$, where $\tau_{\min}$ is defined by (2.5). We now assume that $\tau_{\min} = 0$ and thus y_τ is defined on Ω for all $\tau \geq 0$. This assumption holds if m is large enough and the properties of the algorithm guarantee finite termination for all x in Ω. For an overview of the convergence theory of Krylov solvers, see [15, 16]. For recent results on bicgstab and idr, see [9].

Other properties of y_τ under the assumption that $\tau_{\min} = 0$ follow from Theorem 2.2. We now know that y_τ is continuously differentiable in Ω for almost all τ and that y'_τ is uniformly bounded on Ω for almost all τ. We also note that

$$\lim_{\tau \to 0} y_\tau(x) = A^{-1}h(x), \tag{2.6}$$

since y_τ satisfies the termination test (2.1) for $\tau > 0$ and $A^{-1}h(x)$ is the unique solution to the limiting case $Ay = h(x)$.

An important point to keep in mind is that although y'_τ exists except for a finite number of values of τ, the function y_τ may not be continuous at these exceptional values of τ. The following result illustrates this point.

Theorem 2.3. *The mapping $y_\tau : \mathbb{R}^n \mapsto \mathbb{R}^n$ can be discontinuous if $\tau \in \{\tau_0, \ldots, \tau_m\}$.*

Proof. Consider the conjugate gradient solver for the system $Ay = h(x)$, where

$$A = \begin{pmatrix} \beta & 0 \\ 0 & 1 \end{pmatrix}, \qquad h(x) = \begin{pmatrix} \beta \\ 0 \end{pmatrix}, \qquad \beta > 0.$$

Given a starting point $x_0 = (\xi_1, \xi_2)$, the first iterate of the conjugate gradient method is

$$y_1 = \begin{pmatrix} \xi_1 + \alpha_0\beta(1-\xi_1) \\ (1-\alpha_0)\xi_2 \end{pmatrix}, \quad \text{where} \quad \alpha_0 = \frac{\beta^2(1-\xi_1)^2 + \xi_2^2}{\beta^3(1-\xi_1)^2 + \xi_2^2}.$$

Thus the step length α_0 is a rational function of (ξ_1, ξ_2). Moreover, a calculation shows that

$$\phi(\xi_1, \xi_2) = \|Ay_1 - b\|^2 = \beta^2(1-\xi_1)^2(\alpha_0\beta - 1)^2 + (1-\alpha_0)^2\xi_2^2.$$

The square of the residual, ϕ, and the mapping y_1 are continuously differentiable functions of the initial iterate (ξ_1, ξ_2) in any set that does not contain the solution $y^* = (1, 0)$ of the linear system. If the starting point is chosen in a compact set that does not contain y^*, then there is a $\sigma > 0$ such that $\|y_1 - y^*\| \geq \sigma$.

Restrict the initial point (ξ_1, ξ_2) so that $\nabla\phi(\xi_1, \xi_2) \neq 0$. We now show that y_τ is not a continuous function of the starting point (ξ_1, ξ_2) for $\tau = \tau_1$. Note that $y_{\tau_1}(\xi_1, \xi_2) = y_1$ by the definition of τ_1. Since ϕ increases along any ray p with $\nabla\phi(\xi_1, \xi_2)^T p > 0$, a small perturbation to the initial iterate along the ray p increases the residual, and then the conjugate gradient method takes one more iteration to achieve the tolerance τ_1. The finite termination property of the conjugate gradient method then guarantees that $y_{\tau_1}(\xi_1, \xi_2) = y^*$ for the perturbed initial point. Since $\|y_1 - y^*\| \geq \sigma$, we have shown that a small change in the starting point can yield changes that are bounded away from zero. Hence, y_{τ_1} is not continuous. □

We used the conjugate gradient method to illustrate that y_τ may be discontinuous, but Theorem 2.3 can be modified for other Krylov methods. The important characteristics are that the square of the residual is a rational function of the initial iterate and that the solver terminates at the solution after two iterates when $n = 2$.

3 Noise Levels of Krylov Solvers

We now investigate the noise level of a function defined by the solution of a system of linear equations by an iterative process. The model function that we have chosen is defined by

$$f_\tau(x) = \|y_\tau(x)\|^2, \tag{3.1}$$

where $y_\tau(x)$ is the first iterate of a Krylov solver that satisfies the termination test (2.1). Section 2 presents results on the continuity and differentiability of y_τ for a general system of the form $Ay = h(x)$; our computational results use $h(x) = x$ exclusively.

We consider a family of functions f_τ defined by a matrix A and indexed by the tolerance τ. 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 [6]. Following [13], we scale the matrices by their diagonals and randomly select the base point x_0.

Almost 25% of the matrices in this selection have a small condition number. Fifteen of the matrices are diagonal and hence have unit condition number after the scaling; 28 matrices have condition number less than 10; and all but 10 have condition number less than 10^{10}.

We begin our study of the properties of f_τ by computing the noise level. We can estimate the noise level by evaluating f_τ in a neighborhood and calculating the standard deviation of the function values. This approach is not computationally feasible for expensive functions, however, and thus we have developed the ECNoise algorithm [13] to determine the noise level ε_f of a function in a few function evaluations.

The noise level ε_f provides valuable information on the function in a neighborhood of the base point x_0. For example, we can use the noise level ε_f of a function f in a neighborhood of x_0 to make assertions about other function values in this neighborhood. In particular, we can claim that if we make small perturbations to x_0, then

$$|f(x) - f(x_0)| \leq \gamma\varepsilon_f$$

is likely to hold for $\gamma \geq 1$ of modest size. The noise level ε_f models the uncertainty in the function because, as discussed in [13], small changes in the code that evaluates f (for example, the order of operations, compiler options, libraries, or operating system) produce average variations of order ε_f in the value of f. Thus, the number of significant digits in $f(x)$ is likely to be less than $\lfloor \log(\varepsilon_f^{-1}) \rfloor$.

We have emphasized the role that the noise level ε_f plays in the uncertainty of the function values, but the noise level also plays an important role in termination criteria for optimization algorithms and in determining optimal difference parameters. See [10, 12, 13] for a discussion of these issues and more information on computational noise. Sections 4 and 5 discuss, respectively, the relationship between the noise level ε_f and the accuracy of derivatives and differences.

The ECNoise algorithm [13] estimates the noise level ε_f of a function by evaluating f along a ray in the domain of the function. Specifically, ECNoise chooses a random direction p and evaluates the function $t \mapsto f_\tau(x_0 + tp)$ on a set of n_f equally spaced points on the interval $[-h, h]$ for some parameter h. As shown in [12, 13], the performance of ECNoise is essentially independent of the choices of h and p provided h is reasonably small.

The ECNoise algorithm determines the absolute noise in a neighborhood of a point x_0, but in discussions of the noise level we find it convenient to use ε_f to refer to the relative noise of f_τ, that is, the absolute noise divided by (nonzero) $|f(x_0)|$. We expect the relative noise to be on the order of machine precision for the mathematical functions that are typically provided by a programming environment, but the noise level of more complicated functions must be determined computationally.

From the benchmark set of 116 matrices we looked at f_τ for a variety of tolerances and

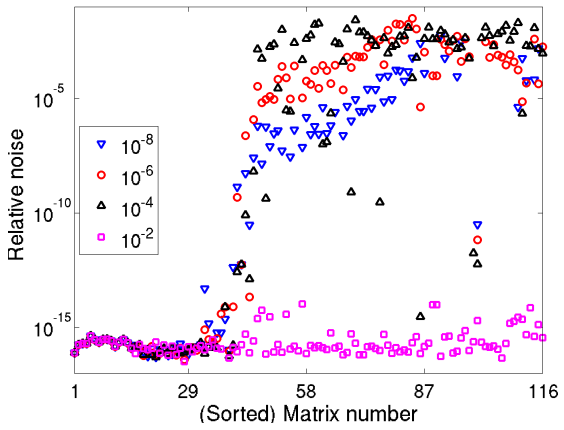

Figure 3.1: Distribution of the relative noise ε_f for f_τ using bicgstab.

solvers. Here we focus on data for the bicgstab solver and tolerances in

$$\mathcal{T} = \{10^{-k} : k = 2, 4, 6, 8\}. \tag{3.2}$$

Results for more tolerances and solvers can be found at [11].

For bicgstab, convergence (satisfying (2.1) within n iterations, where n is the size of x_0) was obtained in more than 90% of these 464 problems, with failures usually occurring when the associated matrix was ill-conditioned and the tolerance τ was small. We now restrict attention to the 435 cases where the Krylov solver converges and thus f_τ is well defined. In these cases ECNoise with $n_f = 9$ and $h = 10^{-12}$ computed the noise level in more than 99% of the problems. ECNoise is not designed for very noisy problems and does not return an estimate when the relative noise is much larger than 10^{-2}; the largest relative noise returned on these problems was approximately $3 \cdot 10^{-2}$.

Figure 3.1 shows the relative noise ε_f for the functions f_τ defined by (3.1) with the four tolerances $\tau \in \mathcal{T}$. In this plot we have arranged the functions in increasing order of the condition number of the matrix that defines the function. Thus, functions with well-conditioned matrices are listed first. This plot shows that there is a set of roughly 35 functions with a noise level on the order of machine precision for all tolerances. This result agrees with our expectations because, for these problems, the Krylov solver converges in a few iterations, often independent of the tolerance τ.

The most interesting functions are those for which the noise level is at least an order of magnitude above machine precision for at least one tolerance. These functions are outside the first third of Figure 3.1. In general, at least 80% of these functions have a noise level above 10^{-10} for at least one tolerance. This percentage depends on the solver; it goes to 90% for bicgstab, but declines to 71% for minresqlp [5] and 47% for gmres. These observations are of interest because they imply that, with high probability, most solvers deliver at most 10 significant digits in the value of f_τ for some tolerance, and this places a limit on the precision of any simulation that uses this value. By this measure, gmres is the solver that generates the least noise.

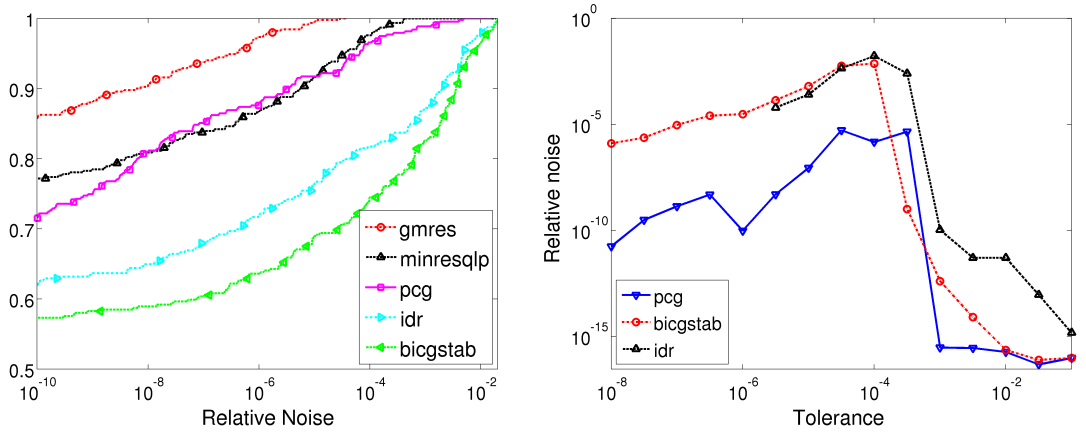

Figure 3.2: Distribution of the relative noise for f_τ. On the left is the cumulative distribution as a function of ε_f for five different Krylov solvers from [11]. On the right is the distribution of ε_f as a function of τ for f_τ using a matrix with the median condition number.

We can obtain an accurate view of the noise generated by different solvers by computing the cumulative distribution function for ε_f generated by a solver. Figure 3.2 (left) shows the portion of the empirical distribution of the relative noise in f_τ over the range $[10^{-10}, 10^{-2}]$. Thus, gmres generates the least noise, idr (idr(s) with $s = 4$ is used throughout our experiments) and bicgstab generate the most noise, while minresqlp generates about the same amount of noise as pcg. Additional information can be obtained from the histograms of the relative noise in [11].

Another interesting observation is that the noise level decreases as the tolerance is decreased, but only after the tolerance decreases below a certain threshold that varies from function to function. This can be deduced from Figure 3.1 by noting that the data points for $\tau = 10^{-8}$ are almost always below the data point for $\tau = 10^{-6}$. However, the position of the data point for $\tau = 10^{-4}$ is unpredictable because it tends to be near the threshold where the noise level starts to decrease.

Figure 3.2 (right) underscores this behavior on a typical (median condition number in the benchmark set) matrix. This figure lends further evidence to the consistency of ECNoise, since the noise level varies remarkably smoothly as the tolerance changes; the lack of data points for idr with small tolerances is due to the lack of convergence in n iterations. For the three Krylov solvers bicgstab, pcg, and idr, the noise is on the order of machine precision when the tolerance τ is large but then sharply rises as the number of iterations (and hence the number of operations) increases. When τ is sufficiently small, the noise begins to decline.

The relationship between the noise and the tolerance in Figure 3.2 is typical of that seen for other solvers and other matrices [11]. This behavior is interesting because it illustrates that a nonmonotone relationship exists between the noise and the truncation error (as guided by the tolerance). While decreasing a tolerance may guarantee smaller truncation errors, it can also lead to a significant rise in the noise. As discussed in the next section, high noise levels can adversely impact applications where derivatives are desired.

4 Computational Noise in Derivatives

In this section we examine the accuracy of derivative computations by comparing different methods for computing a derivative. In many cases, the differences are small, and so the relative errors are an adequate measure of the differences. In some cases, however, we compare quantities that differ by several orders of magnitude, and then we need a different measure. The relative error is the standard method to compare scalars that differ in magnitude, but we use a variation that is symmetric with respect to the scalars used in the comparisons. We define $\mathsf{re}\,(\alpha,\beta)$ for scalars α and β with $\max(|\alpha|,|\beta|) > 0$ by

$$\mathsf{re}\,(\alpha,\beta) = \frac{|\alpha-\beta|}{\max(|\alpha|,|\beta|)}, \tag{4.1}$$

and we complete this definition by setting $\mathsf{re}\,(\alpha,\beta) = 0$ when $\max(|\alpha|,|\beta|) = 0$. Ziv [17] proposed this definition for the analysis of roundoff errors and proved that $\mathsf{re}\,(\cdot,\cdot)$ was a metric for $\mathbb{R}$. Below we gather additional properties of this metric.

Lemma 4.1. *Given α and β in $\mathbb{R}$, the function defined in (4.1) satisfies the following:*

- $\mathsf{re}\,(\alpha,\beta) = 0$ *if and only if* $\alpha = \beta$.
- $\mathsf{re}\,(\alpha,\beta) \in (0,1)$ *if and only if α and β have the same sign and* $\alpha \neq \beta$.
- $\mathsf{re}\,(\alpha,\beta) = 1$ *if and only if* $0 = \min(|\alpha|,|\beta|) < \max(|\alpha|,|\beta|)$.
- $\mathsf{re}\,(\alpha,\beta) \in (1,2]$ *if and only if α and β have opposite signs.*

Proof. The proof follows the same pattern in all cases, so we consider only the second case.

Assume first that $\mathsf{re}\,(\alpha,\beta) \in (0,1)$. Then $\alpha \neq \beta$, and neither α nor β can be zero. If, on the contrary, we have α and β with opposite signs, then we can assume that $\alpha > 0 > \beta$. This implies that

$$|\alpha-\beta| > \max(\alpha,-\beta) = \max(|\alpha|,|\beta|),$$

and then $\mathsf{re}\,(\alpha,\beta) > 1$. This contradiction shows that α and β must have the same sign.

Assume now that α and β have the same sign. Without loss of generality we can also assume that $\alpha \geq \beta > 0$. In this situation,

$$|\alpha-\beta| < |\alpha| = \max(|\alpha|,|\beta|).$$

Hence, $\mathsf{re}\,(\alpha,\beta) < 1$. Moreover, if $\alpha \neq \beta$, we also have $\mathsf{re}\,(\alpha,\beta) > 0$. □

Lemma 4.1 highlights some of the advantages of the relative error defined by $\mathsf{re}\,(\cdot,\cdot)$ over the standard relative error $|\alpha-\beta|/|\alpha|$. Of key importance is the symmetry of $\mathsf{re}\,(\cdot,\cdot)$ with respect to the arguments so that the relative error does not depend on the order of the arguments. This definition of relative error can be extended to vectors in $\mathbb{R}^n$ by defining

$$\mathsf{re}\,(x,y) = \frac{\|x-y\|}{\max(\|x\|,\|y\|)}.$$

This extension is of interest because Ziv [17] has shown that if $\|\cdot\|$ is an inner-product norm, then $(x, y) \mapsto \mathsf{re}\,(x, y)$ is a metric. Another possible extension is to functions defined in a domain Ω. In this case,

$$\mathsf{re}\,(f_1, f_2) = \sup_{x\in\Omega} \mathsf{re}\,(f_1(x), f_2(x))\,,$$

where $f_k : \mathbb{R}^n \mapsto \mathbb{R}^m$ for $k = 1, 2$. Since $\mathsf{re}\,(\cdot,\cdot)$ is a metric in $\mathbb{R}^m$, this definition extends $\mathsf{re}\,(\cdot,\cdot)$ to function space.

We use $\mathsf{re}\,(\cdot,\cdot)$ to measure the difference between two data vectors. A plot of the relative error between components of the data vectors leads to nicely scaled graphs because in all cases $\mathsf{re}\,(\alpha, \beta) \leq 2$. In these plots we pay special attention to components with $\mathsf{re}\,(\alpha, \beta) \geq 1$ because Lemma 4.1 shows that these components do not have the same sign. The next result shows that components with $\mathsf{re}\,(\alpha, \beta) \approx 1$ also merit close attention.

Lemma 4.2. *If* $\mathsf{re}\,(\alpha, \beta) = \mu \in (0, 1)$, *then*

$$\max(|\alpha|, |\beta|) = \frac{1}{1-\mu} \min(|\alpha|, |\beta|).$$

Proof. If $\mathsf{re}\,(\alpha, \beta) \in (0, 1)$, then α and β have the same sign. Assume, without loss of generality, that $\alpha \geq \beta > 0$. Then

$$\mu = \mathsf{re}\,(\alpha, \beta) = \frac{\alpha - \beta}{\alpha},$$

and thus $(1 - \mu)\alpha = \beta$, as we wanted to show. □

We use $\mathsf{re}\,(\cdot,\cdot)$ to examine the accuracy of derivative computations by computing the directional derivative $f'_\tau(x_0; p)$ of the functions f_τ defined by (3.1). We use the same random direction p as in the estimation of the noise. For these experiments we use the IntLab [14] and ADiMat [1] systems for computing the directional derivative. The quantity of interest in these experiments is

$$\Gamma_2(f'_\tau) \equiv \mathsf{re}\left(f'_{\tau,1}(x_0; p), f'_{\tau,2}(x_0; p)\right), \tag{4.2}$$

where $f'_{\tau,1}(x_0; p)$ and $f'_{\tau,2}(x_0; p)$ are the directional derivatives computed by the two systems. Thus, $\Gamma_2(f'_\tau)$ is the relative error of f'_τ when computed by two different algorithms. IntLab and ADiMat use the forward mode of automatic differentiation [8] to compute these derivatives, but the implementation of this mode (operator overloading and source transformation, respectively) can lead to different algorithms.

The relative error $\Gamma_2(f'_\tau)$ measures the accuracy of the derivative computed by automatic differentiation tools, but requires two systems for the generation of the derivatives. This is a serious obstacle for many scientific computations. For example, automatic differentiation through operator overloading was not performed for some of the other solvers tested in our experiments because it required significant changes to the standard implementations. As an alternate measure, consider

$$\Gamma(f'_\tau) \equiv \mathsf{re}\left(f'_\tau(x_0; p), f'_\tau(x_0; (1+\varepsilon)p)\right), \qquad \varepsilon > 0,$$

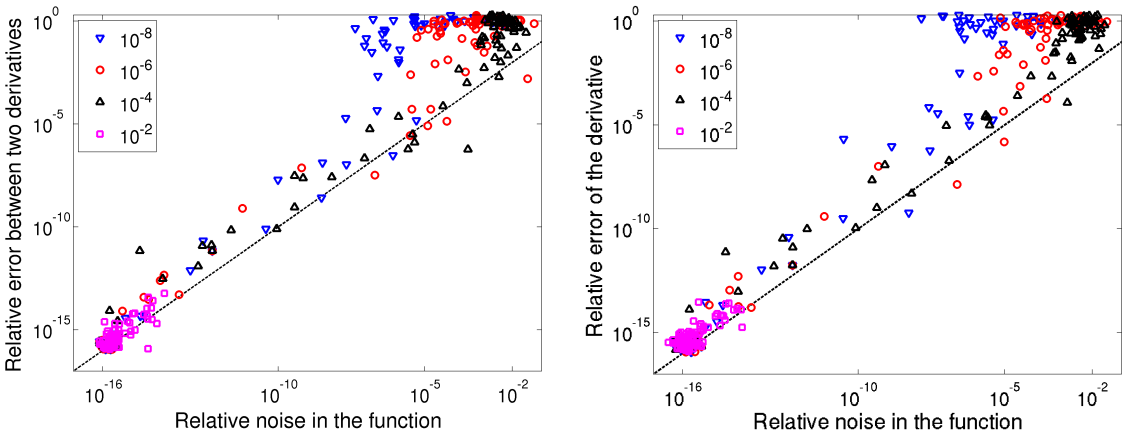

Figure 4.1: Distribution of the relative error estimates $\Gamma_2(f'_\tau)$ (left) and $\Gamma(f'_\tau)$ (right) for bicgstab as a function of the relative noise ε_f; the dashed line is (t, t).

where p is a given direction in $\mathbb{R}^n$. The measure $\Gamma(f'_\tau)$ is the relative error when f'_τ is computed at a fixed x_0 with a small modification to the direction p. For any $\varepsilon > 0$ we have $\Gamma(f'_\tau) = \varepsilon/(1+\varepsilon)$ since the directional derivative is a linear function of the second argument. Thus $\Gamma(f'_\tau)$ is independent of the function f_τ, the base point x_0, and the direction p. Note, however, that although $\Gamma(f'_\tau) \leq \varepsilon$, this may not hold when $f'_\tau(x_0; p)$ is computed in finite precision.

Figure 4.1 compares these two measures of accuracy for the bicgstab solver. On the left side is the distribution of $\Gamma_2(f'_\tau)$ while the right side shows the values of $\Gamma(f'_\tau)$ using the ADiMat derivative and $\varepsilon = \varepsilon_M$, where ε_M is the machine precision in double-precision IEEE arithmetic. This figure shows that the results obtained with $\Gamma_2(f'_\tau)$ and $\Gamma(f'_\tau)$ are similar, and thus either measure could be used to measure accuracy. The measure $\Gamma(f'_\tau)$ is preferable to $\Gamma_2(f'_\tau)$ because it can be computed for any simulation for which we have a derivative code at the cost of one additional evaluation of the directional derivative, and thus we use $\Gamma(f'_\tau)$ with $\varepsilon = \varepsilon_M$ and the ADiMat derivative throughout the rest of the paper.

One of the most interesting aspects of Figure 4.1 is the appearance of high relative errors $\Gamma_2(f'_\tau)$ and $\Gamma(f'_\tau)$ for functions with high levels of noise. Functions with $\Gamma_2(f'_\tau) \geq 0.5$ are of special interest because for these functions the derivatives computed by IntLab and ADiMat differ by at least a factor of 2. A detailed examination of the data shows that for these functions the derivatives often differ in sign and the magnitude grows as τ decreases. Hence, the computational evidence indicates that the directional derivative f'_τ is not bounded as τ decreases. Since Theorem 2.2 guarantees that the derivatives are bounded in infinite precision, these results show that the theoretical properties of the Krylov solver do not imply accurate derivatives in finite precision as the tolerance τ decreases.

Figure 4.1 (right) suggests that the noise level ε_f can be used as an estimate of the relative error $\Gamma(f'_\tau)$ in the derivative on about half the problems and that $\Gamma(f'_\tau)$ is high on most of the other problems. This observation holds for all the solvers in [11] because an analysis of the data shows that if we consider problems with reasonably accurate derivatives

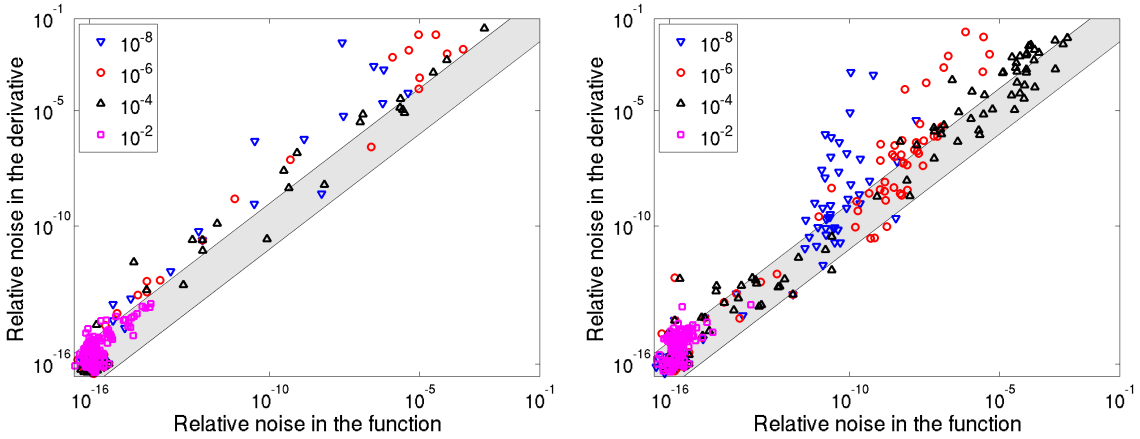

Figure 4.2: Distribution of the noise level in the ADiMat derivatives for bicgstab (left) and pcg (right); the shaded region lies between $(t, t/10)$ and $(t, 10t)$.

(in the sense that $\Gamma(f'_\tau) \leq 0.05$), then the noise level ε_f is within a factor of 100 of $\Gamma(f'_\tau)$ for at least 80% of the problems. This observation shows that we cannot expect to obtain the derivative with double-precision accuracy if the noise level ε_f of the function is safely above machine precision, and thus claims of obtaining derivatives to working accuracy are false.

We have also investigated the relationship between the noise level of the directional derivative and the noise level of the function. For bicgstab, Figure 4.2 (left) shows that roughly 80% of the computable noise estimates are within a factor 10 of the noise in the function, ε_f. As previously discussed, ECNoise does not return an estimate when the noise is too large; if derivatives are at least as noisy as functions, it is not surprising that many data points are missing in Figure 4.2 (left) for the noisiest functions. Figure 4.2 (right) shows the noise in the derivatives for pcg, which tends to be less noisy than bicgstab. Our results [11] show that for all Krylov solvers, the noise level of the derivative is larger than $\varepsilon_f/10$ (the lower boundary of the shaded region in Figure 4.2) on 97% of the problems.

As an outcome of this work, we know that the noise level of the derivative tends to be an order of magnitude estimate of the accuracy as measured by $\Gamma(f'_\tau)$. This is reassuring because obtaining an accurate estimate of the noise level of the derivative tends to require the evaluation of the derivative at six or more nearby points, whereas $\Gamma(f'_\tau)$ requires only one additional derivative evaluation. We also know that the noise level ε_f of the function tends to be an estimate, up to a factor of 100, of the accuracy $\Gamma(f'_\tau)$.

5 Reliable Estimates for Forward Differences

We now examine how the noise level affects the accuracy of a forward difference approximation to the derivative. In earlier work [12] we showed that we can determine a parameter h^* that is provably near-optimal when the function values are determined by a stochastic process. Our numerical results showed that we obtain accurate derivative estimates even

when the function is deterministic; but those results used $\tau = 10^{-3}$ as the tolerance in the iterative solver, and thus the resulting functions f_τ had relatively low noise levels, as illustrated in Figure 3.1. In this section we examine how the accuracy of the derivative estimate depends on the noise level of the function for a wide range of tolerances.

The results of the optimality of a difference parameter [12] depend on the assumption that the function values are determined by a stochastic process. Specifically, assume that $f : \mathbb{R} \mapsto \mathbb{R}$ is defined on an interval I by

$$f(t) = f_s(t) + \varepsilon(t), \qquad t \in I, \tag{5.1}$$

where $f_s : \mathbb{R} \mapsto \mathbb{R}$ is a smooth, deterministic function and the random variables $\varepsilon(t)$ are independent and identically distributed for all $t \in I$. Under these assumptions, the noise level of f is the standard deviation

$$\varepsilon_f = \left(\mathrm{Var}\{\varepsilon(t)\}\right)^{1/2}.$$

An optimal difference parameter $h^* > 0$ minimizes the root mean squared error

$$\mathcal{E}(h) = \mathrm{E}\left\{\left(\frac{f(t_0+h) - f(t_0)}{h} - f_s'(t_0)\right)^2\right\}^{1/2}.$$

One of the main results in [12] is that if $I = [t_0, t_0 + h_0]$ and f_s is twice differentiable, with μ_L and μ_M denoting the minimum and maximum of $|f_s''|$ on I, then

$$(\gamma_1 \mu_L \varepsilon_f)^{1/2} \le \min_{0 \le h \le h_0} \mathcal{E}(h) \le (\gamma_1 \mu_M \varepsilon_f)^{1/2}, \qquad \gamma_1 = 2^{1/2}, \tag{5.2}$$

provided that $h_0 \ge h^*$ for

$$h^* = \gamma_2 \left(\frac{\varepsilon_f}{\mu_M}\right)^{1/2}, \qquad \gamma_2 = 8^{1/4} \approx 1.68. \tag{5.3}$$

The constraint $h_0 \ge h^*$ guarantees that the interval defined by h_0 is not too small. The results in [12] also show that the parameter h^* in (5.3) is a good estimate for the minimizer of $\mathcal{E}$, that is,

$$(\gamma_1 \mu_L \varepsilon_f)^{1/2} \le \mathcal{E}(h^*) \le (\gamma_1 \mu_M \, \varepsilon_f)^{1/2}. \tag{5.4}$$

These bounds are narrow when f_s'' does not change significantly over I and can be used to show that h^* is nearly optimal, because the error obtained using h^* is within a constant factor of the minimal error.

The theory that we have outlined assumes that the function values are determined by a stochastic process, but these results indicate how to determine a nearly optimal difference parameter h when f is deterministic. All that is needed are estimates of the noise level ε_f and of $|f''|$ near t_0; then (5.3) can be used to determine h with μ_M set to the estimate of $|f''|$. We have already discussed the use of ECNoise [13] to determine the noise level; we refer to [12, Section 5] for the details of the $|f''|$ estimator that is used to determine h^* in our results.

We have been discussing an algorithm for functions defined on $\mathbb{R}$, but this algorithm extends to the n-dimensional case, where $f : \mathbb{R}^n \mapsto \mathbb{R}$, by using the algorithm on the function $t \mapsto f(x_0 + tp)$ with $t_0 = 0$. Thus, the algorithm in [12] can be used to determine a difference approximation

$$\delta f(x_0; p) = \frac{f(x_0 + hp) - f(x_0)}{h} \tag{5.5}$$

to the directional derivative $f'(x_0; p)$. If the choice of h is a reasonable approximation to the optimal difference parameter h^* defined by (5.3) and $\mu \approx |f''(t_0)|$, then the bounds (5.4) on the expected error show that

$$\Gamma(\delta f) \equiv \frac{(\gamma_1 \mu \, \varepsilon_f)^{1/2}}{|f'(t_0)|} \tag{5.6}$$

is an estimate of the relative error between $\delta f(x_0; p)$ and $f'(x_0; p)$. The computational results in this section will show that the estimate $\Gamma(\delta f_\tau)$ for the function f_τ defined in (3.1) is an excellent estimate for the relative error

$$\mathsf{re}(\delta f_\tau, f'_\tau) = \mathsf{re}\Big(\delta f_\tau(x_0; p), f'_\tau(x_0; p)\Big) \tag{5.7}$$

between $\delta f_\tau(x_0; p)$ and $f'_\tau(x_0; p)$.

The estimate $\Gamma(\delta f_\tau)$ for the relative error in δf and the estimate $\Gamma(f'_\tau)$ for the relative error in the derivative of f_τ rely on different assumptions. The estimate $\Gamma(\delta f_\tau)$ is based on the assumption that the function satisfies a stochastic model, so it is not strictly valid for deterministic functions. Also, a stochastic bound does not imply a deterministic bound. On the other hand, the estimate $\Gamma(f'_\tau)$ is based on the relative error between directional derivatives along two nearby points in p.

We evaluate the accuracy of the forward-difference approximation (5.5) with the same set of problems as in Sections 3 and 4. For each tolerance $\tau \in \mathcal{T}$ we compute the directional derivative $f'(x_0; p)$ with ADiMat [1] and use the algorithm in [12] to determine an approximation $\delta f(x_0; p)$ to the directional derivative. The ECNoise algorithm computes a noise level estimate for 431 problems, and the $|f''|$ estimator has no Krylov failures for 368 of these problems. More Krylov failures occur in the $|f''|$ estimator because the estimator needs function values at points that are of order $\varepsilon_f^{1/4}$ away from x_0, whereas ECNoise can work with points much closer to x_0. The $|f''|$ estimator returns an estimate on 366 of the 368 possible data points; this performance supports the high level of reliability of the $|f''|$ estimator.

We have plotted the distribution of $\mathsf{re}(\delta f_\tau, f'_\tau)$ in Figure 5.1 for the cases where f_τ is computed with the bicgstab and pcg solvers, but similar results hold for other Krylov solvers [11]. This plot shows the dependence of $\mathsf{re}(\delta f_\tau, f'_\tau)$ on the estimate $\Gamma(\delta f_\tau)$ of the relative error between δf_τ and f'_τ. Data is restricted to points that satisfy $\Gamma(f'_\tau) \le 0.05$ and $\Gamma(\delta f_\tau) \le 0.05$, so we can be reasonably sure that both the directional derivative f'_τ and the difference approximation δf_τ have two significant digits. Otherwise, it would not be sensible to compare δf_τ with f'_τ.

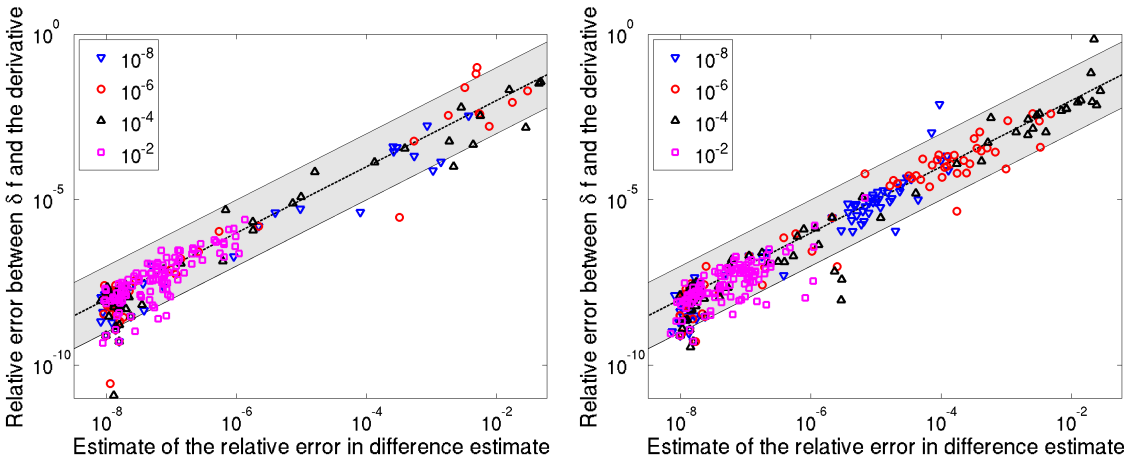

Figure 5.1: Distribution of $\mathsf{re}(\delta f_\tau, f'_\tau)$ for bicgstab (left) and pcg (right) when $\Gamma(f'_\tau) \le 0.05$ and $\Gamma(\delta f_\tau) \le 0.05$. The dashed line is (t, t); the shaded region represents a factor $10^{\pm 1}$ change about the line.

Figure 5.1 shows remarkable agreement between the relative error $\mathsf{re}(\delta f_\tau, f'_\tau)$ and the relative error estimate $\Gamma(\delta f_\tau)$. Since the shaded region represents a factor $10^{\pm 1}$ change about the dashed line (t, t), we conclude that $\Gamma(\delta f_\tau)$ is an order of magnitude estimate of $\mathsf{re}(\delta f_\tau, f'_\tau)$. We emphasize that the data for other Krylov solvers [11] show that this conclusion holds for over 90% of the problems. There is also theoretical support for this conclusion because it can be deduced from (5.4) if we assume that the computed h is close to the theoretical h^* defined by (5.3) and the stochastic bound (5.2) applies in the deterministic case. A few data points fall outside the shaded region, but the deviations are small; in all cases we have $\mathsf{re}(\delta f_\tau, f'_\tau) \le \gamma \Gamma(\delta f_\tau)$ with $\gamma \le 100$.

The computational results in this section suggest that the bound $\Gamma(\delta f)$ is an order of magnitude estimate for the relative error $\mathsf{re}(\delta f, f')$ for a general function f. We can also relate $\mathsf{re}(\delta f, f')$ to the relative noise level by noting that the stochastic bounds (5.2) suggest that in the deterministic case we can expect that

$$\mathsf{re}(\delta f, f') \approx \left(\frac{(\gamma_1 |f''(t_0)|\, |f(t_0)|)^{1/2}}{|f'(t_0)|} \right) \varepsilon_f^{1/2},$$

where ε_f is the relative noise level of f. Under additional assumptions, this estimate shows that $\mathsf{re}(\delta f, f') \approx \varepsilon_f^{1/2}$, a claim found in the literature but usually with little discussion of the underlying assumptions. In this vein we note that the proportionality constant in front of $\varepsilon_f^{1/2}$ depends on f but can be estimated if f exhibits quadratic behavior. Assume, for example, that

$$f(t) = \alpha_0 + \tfrac{1}{2}\alpha_2 (t - t_*)^2,$$

where t_* is the minimizer for f, and that $|\alpha_0|$ is small relative to $f(t_0)$. Under these assumptions, one can easily show that $\mathsf{re}(\delta f, f')$ is of the same order of magnitude as $\varepsilon_f^{1/2}$.

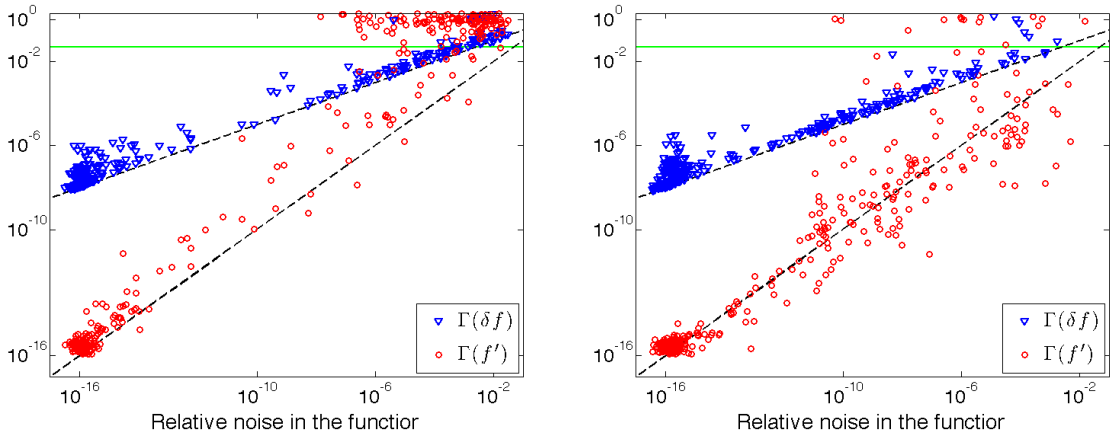

Figure 6.1: Plots of $\Gamma(f'_\tau)$ and $\Gamma(\delta f_\tau)$ for (left) bicgstab and (right) pcg. The lower dashed line is (t, t); the upper line is $(t, t^{1/2})$.

6 Derivatives or Differences?

In the preceding two sections we used the measures $\Gamma(f'_\tau)$ and $\Gamma(\delta f_\tau)$ to analyze the accuracy of derivative and difference approximations as a function of the noise level. We now compare and contrast these findings and thus provide preliminary conclusions about the significance of these results.

For the solvers bicgstab (left) and pcg (right), Figure 6.1 summarizes much of the information that we have obtained by plotting the measures $\Gamma(f'_\tau)$ and $\Gamma(\delta f_\tau)$ for problems where the ECNoise algorithm returns a noise level. Plots for other solvers can be found in [11]. Note that in these plots the data is not restricted to problems where the measures of accuracy suggest that the derivative or the forward difference has two significant digits, that is, when $\Gamma(f'_\tau) \leq 0.05$ or $\Gamma(\delta f_\tau) \leq 0.05$, respectively.

Figure 6.1 shows that if we consider problems with modest levels of noise ($\varepsilon_f \leq 10^{-10}$), then we can reliably compute derivatives or difference approximations with accuracy that generally meets our expectations. Moreover, an analysis of the data in [11] shows that $\Gamma(f'_\tau) \leq 0.05$ with few exceptions, so that we obtain at least two significant digits for at least 98% of the problems. The accuracy of the difference approximation is acceptable with $\Gamma(\delta f_\tau) \leq 0.05$ for at least 96% of the problems and almost all solvers; the exception is gmres, with a 90% percentage.

Figure 6.1 also provides justification for using $\varepsilon_f^{1/2}$ as an order of magnitude estimate of $\Gamma(\delta f_\tau)$ and ε_f as a rough (usually within a factor of 100) estimate of $\Gamma(f'_\tau)$ when the measures of accuracy indicate that the results have at least two significant digits. The plots that appear in [11] show that this observation applies to all the solvers considered.

A striking feature of Figure 6.1 is the large number of data points with $\Gamma(f'_\tau) \geq 0.05$. These points correspond to problems where the derivative does not satisfy minimal accuracy requirements and is often unstable. We emphasize that almost all these points are for problems with high noise levels, that is, $\varepsilon_f > 10^{-10}$. Another interesting feature of Figure 6.1

is that there are few points for which the difference approximation fails to satisfy the accuracy requirement $\Gamma(\delta f_\tau) \leq 0.05$. A partial explanation for this observation is that it is not possible to compute $\Gamma(\delta f_\tau)$ in about 50% of the problems with $\varepsilon_f > 10^{-10}$. As pointed out in Section 5, the algorithm for determining the optimal difference parameter needs to evaluate the function at points that are order $\varepsilon_f^{1/4}$ away from the base point; the Krylov method that determines f_τ tends to fail at these relatively remote points. This situation is not likely to happen in a simulation where we can depend on the function being well-defined when the components of the base point are perturbed by 10%.

A meaningful evaluation of the results for problems with high levels of noise requires that we restrict the data to problems where $\varepsilon_f > 10^{-10}$ and the Krylov solver converges during the estimation of f''_τ. We may still not be able to determine an optimal difference parameter, but this is rare; in four of the solvers the estimation of f''_τ fails on 2% of the problems, and for the other three solvers the estimation fails on 6% of the problems.

The accuracy of the derivative f'_τ for problems with high levels of noise is generally acceptable for all solvers since $\Gamma(f'_\tau) \leq 0.05$ in at least 96% of these problems. However, these percentages drop to 31% and 66% for bicgstab and idr, respectively. The failure to satisfy $\Gamma(f'_\tau) \leq 0.05$ for even 10% of the problems could be considered unacceptable, but the smaller percentages for bicgstab and idr are due to the higher noise levels that are generated by these solvers, as shown in Figure 3.2.

The performance of the difference estimate δf_τ on problems with high levels of noise follows a similar pattern. The percentage of problems with $\Gamma(\delta f_\tau) \leq 0.05$ is at least 90% for most of the solvers, while the solvers with the smallest percentage are bicgstab and idr with 62% and 81%, respectively. An interesting observation on these results is that, contrary to expectations, the results for bicgstab and idr favor the use of differences.

The theoretical and computational results in our work support the conclusion that solvers with high levels of noise are likely to have higher percentage of problems with inaccurate derivatives or difference approximations. Thus, the results for bicgstab and idr could have been predicted from Figure 3.2 by noting that the percentage of problems with high levels of noise level varies significantly between bicgstab and idr and the other solvers.

We have concentrated on the results for bicgstab and idr because these solvers generate the highest level of noise and thus reflect the possible difficulties that may arise with applications with high levels of noise. In this case, as noted above, difference approximations may indeed be more appropriate than derivative approximations. Moreover, as shown in Figure 4.1, the derivatives can become unstable for high levels of noise.

We have shown that it is difficult to reliably compute derivatives or difference approximations for problems with high levels of noise. We may be able to improve the implementation of the function and lower the noise level, but in any case we have provided the tools to quantify the accuracy of these approximations and decide whether the accuracy is suitable for the application under consideration.

Acknowledgments

We are grateful to Daniel Szyld and Sou-Cheng Choi for bringing to our attention new Krylov solvers and their mathematical properties. 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.

References

[1] ADiMat. Automatic Differentiation in Matlab. See http://www.sc.rwth-aachen.de/People/vehreschild/adimat/.

[2] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. V. der Vorst. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods, 2nd ed.* SIAM, Philadelphia, PA, 1994.

[3] T. Beck. Automatic differentiation of iterative processes. *Journal of Computational and Applied Mathematics*, 50(1-3):109–118, 1994. doi:10.1016/0377-0427(94)90293-3.

[4] T. Beck and H. Fischer. The if-problem in automatic differentiation. *Journal of Computational and Applied Mathematics*, 50(1-3):119–131, 1994. doi:10.1016/0377-0427(94)90294-1.

[5] S.-C. T. Choi, C. C. Paige, and M. A. Saunders. MINRES-QLP: A Krylov subspace method for indefinite or singular symmetric systems. *SIAM Journal on Scientific Computing*, 33(4):1810–1836, 2011. doi:10.1137/100787921.

[6] T. A. Davis and Y. Hu. The University of Florida Sparse Matrix Collection. *ACM Transactions on Mathematical Software*, 38:1:1–1:25, 2011. doi:10.1145/2049662.2049663.

[7] J. C. Gilbert. Automatic differentiation and iterative processes. *Optimization Methods and Software*, 1(1):13–21, 1992. doi:10.1080/10556789208805503.

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

[9] M. H. Gutknecht. IDR explained. *Electronic Transactions on Numerical Analysis*, 36(3):126–148, 2010. URL http://etna.ricam.oeaw.ac.at/volumes/2001-2010/vol36/abstract.php?vol=36&pages=126-148.

[10] J. Larson and S. M. Wild. Non-intrusive termination of noisy optimization. *Optimization Methods and Software*, 2012. doi:10.1080/10556788.2012.656116. To appear.

[11] J. J. Moré and S. M. Wild. Do you trust derivatives or differences?: Additional results. See http://www.mcs.anl.gov/~wild/cnoise/.

[12] J. J. Moré and S. M. Wild. Estimating derivatives of noisy simulations. Preprint ANL/MCS-P1785-0810, Argonne National Laboratory, Argonne, Illinois, 2010. To appear in the ACM Transactions on Mathematical Software.

[13] J. J. Moré and S. M. Wild. Estimating computational noise. *SIAM Journal on Scientific Computing*, 33(3):1292–1314, 2011. doi:10.1137/100786125.

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

[15] Y. Saad. *Iterative Methods for Sparse Linear Systems*. SIAM, Philadelphia, PA, 2003.

[16] H. A. van der Vorst. *Iterative Krylov Methods for Large Linear Systems*. Cambridge University Press, 2003.

[17] A. Ziv. Relative distance–an error measure in round-off error analysis. *Mathematics of Computation*, 39(160):563–569, 1982. doi:10.1090/S0025-5718-1982-0669649-2.