Probabilistic Error Analysis for Inner Products

Ilse C.F. Ipsen

Joint work with: Hua Zhou (UCLA)

North Carolina State University
Raleigh, NC, USA

Research supported in part by NSF DMS, and EP/P020720/1 Manchester

Why Statistical/Probabilistic Approaches to Roundoff Error Analysis?

Disadvantage of deterministic bounds:

- Too pessimistic, especially for large dimensions n
  (worst-case bounds cannot account for cancellation of errors)
- Valid only for sufficiently small n
  ($n < \frac{1}{u}$, where u is unit roundoff, half precision: $n < 2048$)
- May specify only first-order error terms

Existing Work

- Von Neumann & Goldstine (1947): Matrix inversion
- Hull & Swenson (1966): Matrix addition, multiplication, Runge Kutta
- Henrici (1966): ODEs
- Tienari (1970): Matrix inversion
- Barlow & Bareiss (1985): Gaussian elimination
- Calvetti (1991, 1992): Convolution, FFT
- Chatelin & Brunet (1990): Eigenvalues
- Higham & Mary (2018):

  Backward errors for: Inner products, matvec, matmult, LU, Cholesky

Overview

1. Perturbation bounds (perturbed inputs, exact computation)
   - General deterministic worst-case bound
   - Probabilistic bound: Independent errors
2. Roundoff error bounds (exact inputs, roundoff in computation)
   - Probabilistic bound: Dependent errors

Assume: All vectors are real

Perturbation Bounds

Perturbed inputs, exact computation

Perturbed Inner Product

Exact vectors

$$\boldsymbol{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \quad \boldsymbol{y} = \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix} \qquad \boldsymbol{x}^T\boldsymbol{y} = x_1 y_1 + \cdots + x_n y_n$$

Perturbed vectors

$$\hat{\boldsymbol{x}} = \begin{pmatrix} (1+\delta_1)x_1 \\ \vdots \\ (1+\delta_n)x_n \end{pmatrix} \quad \hat{\boldsymbol{y}} = \begin{pmatrix} (1+\theta_1)y_1 \\ \vdots \\ (1+\theta_n)y_n \end{pmatrix} \qquad |\delta_j|, |\theta_j| \leq \boldsymbol{u}$$

Relative error

$$\left| \frac{\hat{\boldsymbol{x}}^T\hat{\boldsymbol{y}} - \boldsymbol{x}^T\boldsymbol{y}}{\boldsymbol{x}^T\boldsymbol{y}} \right| \leq \; ??$$

General Deterministic Worst-Case Bound

Idea: Express perturbations as Hadamard product

$$\hat{\boldsymbol{x}} = \begin{pmatrix} x_1 + \delta_1 x_1 \\ \vdots \\ x_n + \delta_n x_n \end{pmatrix} = \boldsymbol{x} + \boldsymbol{\delta} \circ \boldsymbol{x}, \qquad \hat{\boldsymbol{y}} = \begin{pmatrix} y_1 + \theta_1 y_1 \\ \vdots \\ y_n + \theta_n y_n \end{pmatrix} = \boldsymbol{y} + \boldsymbol{\theta} \circ \boldsymbol{y}$$

Relative error bound

$$\left| \frac{\hat{\boldsymbol{x}}^T \hat{\boldsymbol{y}} - \boldsymbol{x}^T \boldsymbol{y}}{\boldsymbol{x}^T \boldsymbol{y}} \right| \leq \underbrace{\frac{\|\boldsymbol{x} \circ \boldsymbol{y}\|_p}{|\boldsymbol{x}^T \boldsymbol{y}|}}_{\text{Amplifier}} \underbrace{\|\boldsymbol{\delta} + \boldsymbol{\theta} + \boldsymbol{\delta} \circ \boldsymbol{\theta}\|_q}_{\text{Perturbations}} \qquad \frac{1}{p} + \frac{1}{q} = 1$$

- $\boldsymbol{\delta} \circ \boldsymbol{\theta}$ represents second-order errors
- Bound is exact

Deterministic Worst-Case Bound: Special Cases

- $p = 1$: Traditional amplifier

$$\left|\frac{\hat{\boldsymbol{x}}^T\hat{\boldsymbol{y}} - \boldsymbol{x}^T\boldsymbol{y}}{\boldsymbol{x}^T\boldsymbol{y}}\right| \leq \frac{|\boldsymbol{x}|^T|\boldsymbol{y}|}{|\boldsymbol{x}^T\boldsymbol{y}|}\, u(2+u)$$

- $p = 2$:

$$\left|\frac{\hat{\boldsymbol{x}}^T\hat{\boldsymbol{y}} - \boldsymbol{x}^T\boldsymbol{y}}{\boldsymbol{x}^T\boldsymbol{y}}\right| \leq \frac{\sqrt{\sum_{j=1}^n |x_j y_j|^2}}{|\boldsymbol{x}^T\boldsymbol{y}|}\, \sqrt{n}\, u(2+u)$$

- $p = \infty$: Smallest amplifier

$$\left|\frac{\hat{\boldsymbol{x}}^T\hat{\boldsymbol{y}} - \boldsymbol{x}^T\boldsymbol{y}}{\boldsymbol{x}^T\boldsymbol{y}}\right| \leq \frac{\max_{1\leq j\leq n} |x_j y_j|}{|\boldsymbol{x}^T\boldsymbol{y}|}\, n\, u(2+u)$$

Comparison of Deterministic Bounds for $n \leq 10^8$

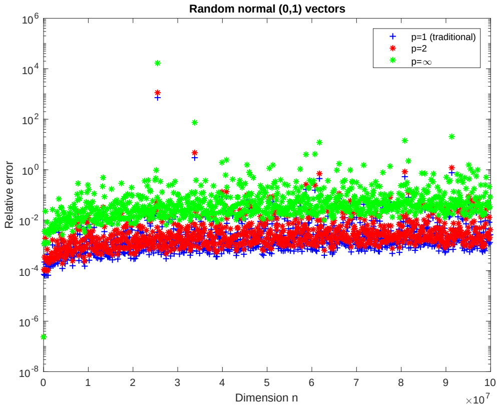

$p = 1$ (traditional) bound is the best
$p = 2$ bound is almost as good

Single precision perturbations $u \approx 10^{-8}$, bounds computed in double

Probabilistic Bound: Azuma's Inequality

How much does a sum $Z = Z_1 + \cdots + Z_n$
of independent random variables $Z_1, \ldots, Z_n$
differ from its mean $\mathbb{E}[Z]$?

If

$$|Z_j - \mathbb{E}[Z_j]| \leq c_j, \qquad 1 \leq j \leq n$$

then

$$\Pr\left[|Z - \mathbb{E}[Z]| \geq \epsilon\right] \leq 2 \exp\left(-\frac{\epsilon^2}{2 \sum_{j=1}^{n} c_j^2}\right)$$

- $2 \sum_{j=1}^{n} c_j^2$ approximates the variance
- All Z_j close to their means $\Rightarrow$
  $Z_1 + \cdots + Z_n$ close to its mean, with high probability

Probabilistic Perturbation Bound

Assume

- All δ_j, θ_j are independent random variables
- Zero mean: $\mathbb{E}[\delta_j] = 0 = \mathbb{E}[\theta_j]$
- Bounded: $|\delta_j|, |\theta_j| \leq u$

Then, for any $0 < \delta < 1$, with probability at least $1 - \delta$

$$\left| \frac{\hat{\boldsymbol{x}}^T \hat{\boldsymbol{y}} - \boldsymbol{x}^T \boldsymbol{y}}{\boldsymbol{x}^T \boldsymbol{y}} \right| \leq \underbrace{\frac{\sqrt{\sum_{j=1}^{n} |x_j y_j|^2}}{|\boldsymbol{x}^T \boldsymbol{y}|}}_{\text{Amplifier}} \underbrace{\sqrt{2 \ln(2/\delta)}}_{\text{Probabilistic}} \underbrace{u(2+u)}_{\text{Perturbation}}$$

- Probabilistic factor is small:
  If $1 - \delta = 1 - 10^{-16}$ then $\sqrt{2 \ln(2/\delta)} \leq 9$

Comparison of Perturbation Bounds: Deterministic vs Probabilistic

$$\left| \frac{\hat{\boldsymbol{x}}^T \hat{\boldsymbol{y}} - \boldsymbol{x}^T \boldsymbol{y}}{\boldsymbol{x}^T \boldsymbol{y}} \right| \leq \frac{\sqrt{\sum_{j=1}^n |x_j y_j|^2}}{|\boldsymbol{x}^T \boldsymbol{y}|} \, \Delta \, u(2+u)$$

- Deterministic ($p = 2$) bound: $\Delta = \sqrt{n}$

  Increases with dimension n

- Probabilistic bound: $\Delta = \sqrt{2 \ln (2/\delta)}$

  Independent of dimension

  $\Delta \leq 9$ for tiny failure probability $\delta = 10^{-16}$

- Probabilistic bound tighter for $n \geq 81$

Comparison: Deterministic vs Probabilistic for $n \leq 10^7$

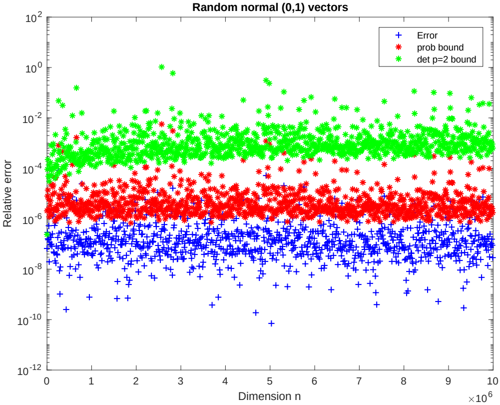

Probabilistic bound tighter than deterministic bound

Single precision perturbations $u \approx 10^{-8}$, bounds computed in double

Comparison: Deterministic vs Probabilistic for $n \le 10^7$

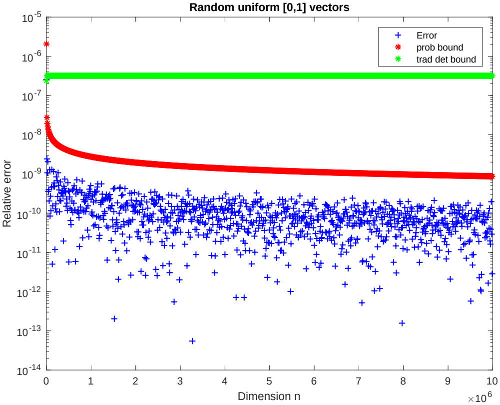

Probabilistic bound tighter than deterministic bound

Single precision perturbations $u \approx 10^{-8}$, bounds computed in double

Perturbation Bounds: Summary

- Component-wise relative perturbation of input vectors, inner product computation is exact
- Bounds for the relative error of $\boldsymbol{x}^T \boldsymbol{y}$
- Deterministic bounds are exact (no big O terms)
  Amplifier in any p-norm
- Probabilistic bound: Perturbations are random variables
- No assumptions on random variables other than:

  independent, zero-mean, bounded

- Probabilistic bound (with stringent success probability) tighter than deterministic bound for dimension $n \geq 81$

Roundoff Error Bounds

Exact inputs, computations have round off errors

Deterministic Roundoff Error Bound

- Exact computation

$$\begin{aligned} s_1 &= x_1 y_1 \\ s_{k+1} &= s_k + x_{k+1} y_{k+1} \qquad 2 \leq k < n \end{aligned}$$

  Output: $s_n = \boldsymbol{x}^T \boldsymbol{y}$

- Floating point arithmetic: $|\delta_k|, |\theta_k| \leq u$

$$\begin{aligned} \hat{s}_1 &= x_1 y_1 (1 + \theta_1) \\ \hat{s}_k &= (\hat{s}_{k-1} + x_k y_k (1 + \theta_k))(1 + \delta_k) \quad 2 \leq k \leq n \end{aligned}$$

  Output: $\hat{s}_n$

- If $nu < 1$ then [Higham 2002]

$$\left| \frac{\hat{s}_n - s_n}{s_n} \right| \leq \underbrace{\frac{|\boldsymbol{x}|^T |\boldsymbol{y}|}{|\boldsymbol{x}^T \boldsymbol{y}|}}_{\text{Amplifier}} \underbrace{\frac{nu}{1 - nu}}_{\text{Roundoff}}$$

Probabilistic Roundoff Error Bound

Distinguish product roundoffs from summation roundoffs

- Exact computation

$$\begin{aligned} s_1 = s_2 &= x_1 y_1 \\ s_{2k+1} = s_{2(k+1)} &= s_{2k} + x_{k+1} y_{k+1} \qquad 2 \leq k < n \end{aligned}$$

Output: $s_{2n} = \boldsymbol{x}^T \boldsymbol{y}$

- Floating point arithmetic: $|\delta_k| \leq u$

$$\begin{aligned} \hat{s}_1 &= x_1 y_1 (1 + \delta_1) \\ \hat{s}_2 &= \hat{s}_1 (1 + \delta_2) \\ \hat{s}_{2k+1} &= \hat{s}_{2k} + x_{k+1} y_{k+1} (1 + \delta_{2k+1}) \\ \hat{s}_{2(k+1)} &= \hat{s}_{2k+1} (1 + \delta_{2(k+1)}) \qquad 2 \leq k < n \end{aligned}$$

Output: $\hat{s}_{2n}$

Probabilistic Bounds for Forward Error

- Product roundoff

$$\begin{aligned} Z_{2k+1} &= \hat{s}_{2k+1} - s_{2k+1} \\ &= Z_{2k} + x_{k+1} y_{k+1}\, \delta_{2k+1} \end{aligned}$$

- Summation roundoff

$$\begin{aligned} Z_{2(k+1)} &= \hat{s}_{2(k+1)} - s_{2(k+1)} \\ &= Z_{2k+1} + \hat{s}_{2k+1} \delta_{2(k+1)} \end{aligned}$$

- Assume that roundoffs δ_k have zero mean: $\mathbb{E}[\delta_k] = 0$

Forward error at stage j, conditioned on previous roundoffs, has mean equal to forward error at stage $j-1$

$$\mathbb{E}[Z_j \,|\, \delta_1, \ldots, \delta_{j-1}] = Z_{j-1} \qquad 1 < j \leq 2n$$

Forward errors $Z_1, Z_2, \ldots,$ are Martingale with respect to roundoffs $\delta_1, \delta_2, \ldots$

Probabilistic Bound: Azuma-Hoeffding Martingale

Sequence of random variables $Z_0, Z_1, Z_2 \ldots$ is Martingale with respect to sequence $\delta_1, \delta_2 \ldots$ if for $j \geq 1$

1. Z_j is function of $\delta_1, \ldots, \delta_j$
2. $\mathbb{E}[|Z_j|] < \infty$,
3. $\mathbb{E}\left[Z_j | \delta_1, \ldots, \delta_{j-1}\right] = Z_{j-1}$

If also

$$|Z_j - Z_{j-1}| \leq c_j \qquad 1 \leq j \leq 2n$$

then, for any $0 < \delta < 1$, with probability at least $1 - \delta$

$$|Z_{2n} - Z_0| \leq \sqrt{\sum_{j=1}^{2n} c_j^2}\,\sqrt{2 \ln\left(2/\delta\right)}$$

Probabilistic Roundoff Bound

Assume that the roundoffs δ_j satisfy:

- Zero mean: $\mathbb{E}[\delta_j] = 0$
- Bounded: $|\delta_j| \leq u$

Then, for any $0 < \delta < 1$, with probability at least $1 - \delta$

$$\left|\frac{\hat{s}_{2n} - s_{2n}}{s_{2n}}\right| \leq \frac{|\mathbf{x}|^T|\mathbf{y}|}{|\mathbf{x}^T\mathbf{y}|}\sqrt{n+1}\ \sqrt{2\ \ln(2/\delta)}\,(1+u)^n\,u$$

Bound does not depend on the summation order

Comparison of Roundoff Bounds: Deterministic vs Probabilistic

$$\left|\frac{\hat{s}_{2n} - s_{2n}}{s_{2n}}\right| \leq \frac{|\mathbf{x}|^T|\mathbf{y}|}{|\mathbf{x}^T\mathbf{y}|}\,\Delta\,u$$

Assume: $\delta = 10^{-16}$, $u \approx 6 \cdot 10^{-8}$ (IEEE Single), $n \leq 10^7$

- Deterministic bound:

$$\Delta = \frac{n}{1 - nu} \leq 1.5\,n$$

- Probabilistic bound:

$$\Delta = \sqrt{n+1}\,(1+u)^n\,\sqrt{2\,\ln(2/\delta)} \leq 15.7\,\sqrt{n+1}$$

- Deterministic bound $\sim n$
  probabilistic bound $\sim \sqrt{n}$

Comparison: Deterministic vs Probabilistic for $n \leq 10^7$

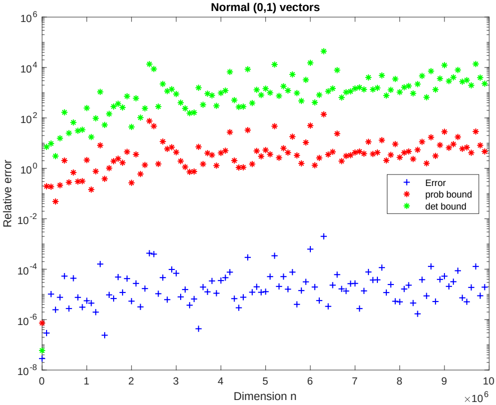

Probabilistic bound tighter than deterministic bound

Single precision perturbations $u \approx 6 \cdot 10^{-8}$, bounds computed in double

Comparison: Deterministic vs Probabilistic for $n \leq 10^7$

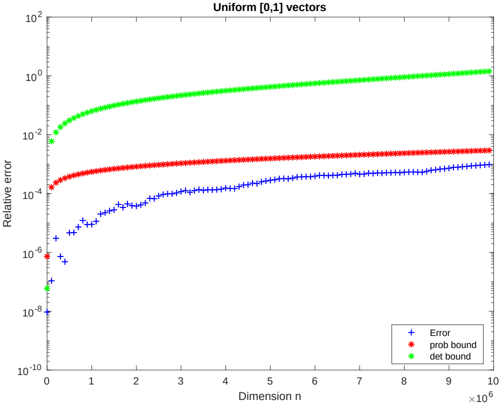

Probabilistic bound tighter than deterministic bound

Single precision perturbations $u \approx 6 \cdot 10^{-8}$, bounds computed in double

Comparison: Deterministic vs Probabilistic for $n \leq 10^8$

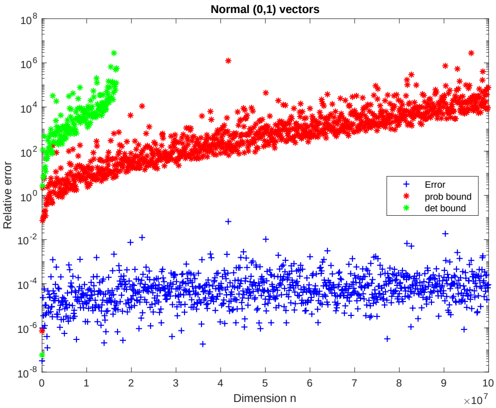

Probabilistic bound tighter than deterministic bound

But: Need to tighten probabilistic bound

Tighter Probabilistic Bound

Assume that the roundoffs δ_j satisfy:

- Zero mean: $\mathbb{E}[\delta_j] = 0$
- Bounded: $|\delta_j| \leq u$

For any $0 < \delta < 1$, with probability at least $1 - \delta$

$$\left| \frac{\hat{s}_{2n} - s_{2n}}{s_{2n}} \right| \leq \kappa \sqrt{2 \ln (2/\delta)}\, u$$

where

$$\kappa \equiv \frac{\sqrt{\|\boldsymbol{x} \circ \boldsymbol{y}\|_2^2 + \sum_{k=1}^{n} \left((\boldsymbol{x} \circ \boldsymbol{y})_{1:k}^T \boldsymbol{u}_k \right)^2}}{|\boldsymbol{x}^T \boldsymbol{y}|}$$

and

$$\boldsymbol{u}_k \equiv \begin{pmatrix} 1 + u & (1+u)^2 & \cdots & (1+u)^k \end{pmatrix}$$

Summary

Probabilistic perturbation bounds:

Relative error independent of n

Probabilistic roundoff error bounds:

- Forward error proportional to $\sim \sqrt{n}$ instead of n
- No limit on dimension n
- No assumption on independence of errors (Martingales)
  Only assumption: zero-mean and bounded
- Exact, non-asymptotic bounds (no big O terms)
- Extremely stringent success probabilities ($\delta = 10^{-16}$)

Not covered:

New condition numbers for general forward error bounds
(from concentration inequalities)