

# NAG Library Function Document

## nag\_dgesvd (f08kbc)

### 1 Purpose

nag\_dgesvd (f08kbc) computes the singular value decomposition (SVD) of a real  $m$  by  $n$  matrix  $A$ , optionally computing the left and/or right singular vectors.

### 2 Specification

```
#include <nag.h>
#include <nagf08.h>

void nag_dgesvd (Nag_OrderType order, Nag_ComputeUType jobu,
                Nag_ComputeVType jobvt, Integer m, Integer n, double a[], Integer pda,
                double s[], double u[], Integer pdu, double vt[], Integer pdvt,
                double work[], NagError *fail)
```

### 3 Description

The SVD is written as

$$A = U\Sigma V^T,$$

where  $\Sigma$  is an  $m$  by  $n$  matrix which is zero except for its  $\min(m, n)$  diagonal elements,  $U$  is an  $m$  by  $m$  orthogonal matrix, and  $V$  is an  $n$  by  $n$  orthogonal matrix. The diagonal elements of  $\Sigma$  are the singular values of  $A$ ; they are real and non-negative, and are returned in descending order. The first  $\min(m, n)$  columns of  $U$  and  $V$  are the left and right singular vectors of  $A$ .

Note that the function returns  $V^T$ , not  $V$ .

### 4 References

Anderson E, Bai Z, Bischof C, Blackford S, Demmel J, Dongarra J J, Du Croz J J, Greenbaum A, Hammarling S, McKenney A and Sorensen D (1999) *LAPACK Users' Guide* (3rd Edition) SIAM, Philadelphia <http://www.netlib.org/lapack/lug>

Golub G H and Van Loan C F (1996) *Matrix Computations* (3rd Edition) Johns Hopkins University Press, Baltimore

### 5 Arguments

1: **order** – Nag\_OrderType *Input*

*On entry:* the **order** argument specifies the two-dimensional storage scheme being used, i.e., row-major ordering or column-major ordering. C language defined storage is specified by **order** = Nag\_RowMajor. See Section 3.2.1.3 in the Essential Introduction for a more detailed explanation of the use of this argument.

*Constraint:* **order** = Nag\_RowMajor or Nag\_ColMajor.

2: **jobu** – Nag\_ComputeUType *Input*

*On entry:* specifies options for computing all or part of the matrix  $U$ .

**jobu** = Nag\_AllU

All  $m$  columns of  $U$  are returned in array **u**.

**jobu** = Nag\_SingularVecsU

The first  $\min(m, n)$  columns of  $U$  (the left singular vectors) are returned in the array **u**.

**jobu** = Nag\_Overwrite

The first  $\min(m, n)$  columns of  $U$  (the left singular vectors) are overwritten on the array **a**.

**jobu** = Nag\_NotU

No columns of  $U$  (no left singular vectors) are computed.

*Constraint:* **jobu** = Nag\_AllU, Nag\_SingularVecsU, Nag\_Overwrite or Nag\_NotU.

3: **jobvt** – Nag\_ComputeVTType *Input*

*On entry:* specifies options for computing all or part of the matrix  $V^T$ .

**jobvt** = Nag\_AllVT

All  $n$  rows of  $V^T$  are returned in the array **vt**.

**jobvt** = Nag\_SingularVecsVT

The first  $\min(m, n)$  rows of  $V^T$  (the right singular vectors) are returned in the array **vt**.

**jobvt** = Nag\_OverwriteVT

The first  $\min(m, n)$  rows of  $V^T$  (the right singular vectors) are overwritten on the array **a**.

**jobvt** = Nag\_NotVT

No rows of  $V^T$  (no right singular vectors) are computed.

*Constraints:*

**jobvt** = Nag\_AllVT, Nag\_SingularVecsVT, Nag\_OverwriteVT or Nag\_NotVT;

If **jobu** = Nag\_Overwrite, **jobvt** cannot be Nag\_OverwriteVT.

4: **m** – Integer *Input*

*On entry:*  $m$ , the number of rows of the matrix  $A$ .

*Constraint:*  $m \geq 0$ .

5: **n** – Integer *Input*

*On entry:*  $n$ , the number of columns of the matrix  $A$ .

*Constraint:*  $n \geq 0$ .

6: **a**[*dim*] – double *Input/Output*

**Note:** the dimension, *dim*, of the array **a** must be at least

$\max(1, \mathbf{pda} \times \mathbf{n})$  when **order** = Nag\_ColMajor;

$\max(1, \mathbf{m} \times \mathbf{pda})$  when **order** = Nag\_RowMajor.

The  $(i, j)$ th element of the matrix  $A$  is stored in

**a**[( $j - 1$ )  $\times$  **pda** +  $i - 1$ ] when **order** = Nag\_ColMajor;

**a**[( $i - 1$ )  $\times$  **pda** +  $j - 1$ ] when **order** = Nag\_RowMajor.

*On entry:* the  $m$  by  $n$  matrix  $A$ .

*On exit:* if **jobu** = Nag\_Overwrite, **a** is overwritten with the first  $\min(m, n)$  columns of  $U$  (the left singular vectors, stored column-wise).

If **jobvt** = Nag\_OverwriteVT, **a** is overwritten with the first  $\min(m, n)$  rows of  $V^T$  (the right singular vectors, stored row-wise).

If **jobu**  $\neq$  Nag\_Overwrite and **jobvt**  $\neq$  Nag\_OverwriteVT, the contents of **a** are destroyed.

- 7: **pda** – Integer *Input*  
*On entry:* the stride separating row or column elements (depending on the value of **order**) in the array **a**.  
*Constraints:*  
 if **order** = Nag\_ColMajor, **pda**  $\geq$   $\max(1, \mathbf{m})$ ;  
 if **order** = Nag\_RowMajor, **pda**  $\geq$   $\max(1, \mathbf{n})$ .
- 8: **s**[*dim*] – double *Output*  
**Note:** the dimension, *dim*, of the array **s** must be at least  $\max(1, \min(\mathbf{m}, \mathbf{n}))$ .  
*On exit:* the singular values of *A*, sorted so that  $\mathbf{s}[i - 1] \geq \mathbf{s}[i]$ .
- 9: **u**[*dim*] – double *Output*  
**Note:** the dimension, *dim*, of the array **u** must be at least  
 $\max(1, \mathbf{pdu} \times \mathbf{m})$  when **jobu** = Nag\_AllU;  
 $\max(1, \mathbf{pdu} \times \min(\mathbf{m}, \mathbf{n}))$  when **jobu** = Nag\_SingularVecsU and **order** = Nag\_ColMajor;  
 $\max(1, \mathbf{m} \times \mathbf{pdu})$  when **jobu** = Nag\_SingularVecsU and **order** = Nag\_RowMajor;  
 $\max(1, \mathbf{m})$  otherwise.  
 The (*i*, *j*)th element of the matrix *U* is stored in  
 $\mathbf{u}[(j - 1) \times \mathbf{pdu} + i - 1]$  when **order** = Nag\_ColMajor;  
 $\mathbf{u}[(i - 1) \times \mathbf{pdu} + j - 1]$  when **order** = Nag\_RowMajor.  
*On exit:* if **jobu** = Nag\_AllU, **u** contains the *m* by *m* orthogonal matrix *U*.  
 If **jobu** = Nag\_SingularVecsU, **u** contains the first  $\min(m, n)$  columns of *U* (the left singular vectors, stored column-wise).  
 If **jobu** = Nag\_NotU or Nag\_Overwrite, **u** is not referenced.
- 10: **pdu** – Integer *Input*  
*On entry:* the stride separating row or column elements (depending on the value of **order**) in the array **u**.  
*Constraints:*  
 if **order** = Nag\_ColMajor,  
   if **jobu** = Nag\_AllU, **pdu**  $\geq$   $\max(1, \mathbf{m})$ ;  
   if **jobu** = Nag\_SingularVecsU, **pdu**  $\geq$   $\max(1, \mathbf{m})$ ;  
   otherwise **pdu**  $\geq$  1.;  
 if **order** = Nag\_RowMajor,  
   if **jobu** = Nag\_AllU, **pdu**  $\geq$   $\max(1, \mathbf{m})$ ;  
   if **jobu** = Nag\_SingularVecsU, **pdu**  $\geq$   $\max(1, \min(\mathbf{m}, \mathbf{n}))$ ;  
   otherwise **pdu**  $\geq$  1..
- 11: **vt**[*dim*] – double *Output*  
**Note:** the dimension, *dim*, of the array **vt** must be at least  
 $\max(1, \mathbf{pdvt} \times \mathbf{n})$  when **jobvt** = Nag\_AllVT;  
 $\max(1, \mathbf{pdvt} \times \mathbf{n})$  when **jobvt** = Nag\_SingularVecsVT and **order** = Nag\_ColMajor;  
 $\max(1, \min(\mathbf{m}, \mathbf{n}) \times \mathbf{pdvt})$  when **jobvt** = Nag\_SingularVecsVT and  
**order** = Nag\_RowMajor;  
 $\max(1, \min(\mathbf{m}, \mathbf{n}))$  otherwise.  
 The (*i*, *j*)th element of the matrix is stored in  
 $\mathbf{vt}[(j - 1) \times \mathbf{pdvt} + i - 1]$  when **order** = Nag\_ColMajor;  
 $\mathbf{vt}[(i - 1) \times \mathbf{pdvt} + j - 1]$  when **order** = Nag\_RowMajor.

On exit: if **jobvt** = Nag\_AllVT, **vt** contains the  $n$  by  $n$  orthogonal matrix  $V^T$ .

If **jobvt** = Nag\_SingularVecsVT, **vt** contains the first  $\min(m, n)$  rows of  $V^T$  (the right singular vectors, stored row-wise).

If **jobvt** = Nag\_NotVT or Nag\_OverwriteVT, **vt** is not referenced.

12: **pdvt** – Integer *Input*

On entry: the stride separating row or column elements (depending on the value of **order**) in the array **vt**.

Constraints:

```

if order = Nag_ColMajor,
    if jobvt = Nag_AllVT, pdvt  $\geq$  max(1, n);
    if jobvt = Nag_SingularVecsVT, pdvt  $\geq$  max(1, min(m, n));
    otherwise pdvt  $\geq$  1.;
if order = Nag_RowMajor,
    if jobvt = Nag_AllVT, pdvt  $\geq$  max(1, n);
    if jobvt = Nag_SingularVecsVT, pdvt  $\geq$  max(1, n);
    otherwise pdvt  $\geq$  1..

```

13: **work**[min(**m**, **n**)] – double *Output*

On exit: if **fail.code** = NE\_CONVERGENCE, **WORK**(2 : min(**m**, **n**)) (using the notation described in Section 3.2.1.4 in the Essential Introduction) contains the unconverged superdiagonal elements of an upper bidiagonal matrix  $B$  whose diagonal is in **s** (not necessarily sorted).  $B$  satisfies  $A = UBV^T$ , so it has the same singular values as  $A$ , and singular vectors related by  $U$  and  $V^T$ .

14: **fail** – NagError \* *Input/Output*

The NAG error argument (see Section 3.6 in the Essential Introduction).

## 6 Error Indicators and Warnings

### NE\_ALLOC\_FAIL

Dynamic memory allocation failed.  
See Section 3.2.1.2 in the Essential Introduction for further information.

### NE\_BAD\_PARAM

On entry, argument  $\langle value \rangle$  had an illegal value.

### NE\_CONVERGENCE

If nag\_dgesvd (f08kbc) did not converge, **fail.errnum** specifies how many superdiagonals of an intermediate bidiagonal form did not converge to zero.

### NE\_ENUM\_INT\_2

On entry, **jobu** =  $\langle value \rangle$ , **pdu** =  $\langle value \rangle$  and **m** =  $\langle value \rangle$ .  
Constraint: if **jobu** = Nag\_AllU, **pdu**  $\geq$  max(1, **m**);  
if **jobu** = Nag\_SingularVecsU, **pdu**  $\geq$  max(1, **m**);  
otherwise **pdu**  $\geq$  1.

On entry, **jobvt** =  $\langle value \rangle$ , **pdvt** =  $\langle value \rangle$ , **n** =  $\langle value \rangle$ .  
Constraint: if **jobvt** = Nag\_AllVT, **pdvt**  $\geq$  max(1, **n**);  
if **jobvt** = Nag\_SingularVecsVT, **pdvt**  $\geq$  max(1, **n**);  
otherwise **pdvt**  $\geq$  1.

**NE\_ENUM\_INT\_3**

On entry, **jobu** =  $\langle value \rangle$ , **pdu** =  $\langle value \rangle$ , **m** =  $\langle value \rangle$  and **n** =  $\langle value \rangle$ .  
 Constraint: if **jobu** = Nag\_AllU, **pdu**  $\geq \max(1, \mathbf{m})$ ;  
 if **jobu** = Nag\_SingularVecsU, **pdu**  $\geq \max(1, \min(\mathbf{m}, \mathbf{n}))$ ;  
 otherwise **pdu**  $\geq 1$ .

On entry, **jobvt** =  $\langle value \rangle$ , **pdvt** =  $\langle value \rangle$ , **m** =  $\langle value \rangle$  and **n** =  $\langle value \rangle$ .  
 Constraint: if **jobvt** = Nag\_AllVT, **pdvt**  $\geq \max(1, \mathbf{n})$ ;  
 if **jobvt** = Nag\_SingularVecsVT, **pdvt**  $\geq \max(1, \min(\mathbf{m}, \mathbf{n}))$ ;  
 otherwise **pdvt**  $\geq 1$ .

**NE\_INT**

On entry, **m** =  $\langle value \rangle$ .  
 Constraint: **m**  $\geq 0$ .

On entry, **n** =  $\langle value \rangle$ .  
 Constraint: **n**  $\geq 0$ .

On entry, **pda** =  $\langle value \rangle$ .  
 Constraint: **pda**  $> 0$ .

On entry, **pdu** =  $\langle value \rangle$ .  
 Constraint: **pdu**  $> 0$ .

On entry, **pdvt** =  $\langle value \rangle$ .  
 Constraint: **pdvt**  $> 0$ .

**NE\_INT\_2**

On entry, **pda** =  $\langle value \rangle$  and **m** =  $\langle value \rangle$ .  
 Constraint: **pda**  $\geq \max(1, \mathbf{m})$ .

On entry, **pda** =  $\langle value \rangle$  and **n** =  $\langle value \rangle$ .  
 Constraint: **pda**  $\geq \max(1, \mathbf{n})$ .

**NE\_INTERNAL\_ERROR**

An internal error has occurred in this function. Check the function call and any array sizes. If the call is correct then please contact NAG for assistance.

An unexpected error has been triggered by this function. Please contact NAG.  
 See Section 3.6.6 in the Essential Introduction for further information.

**NE\_NO\_LICENCE**

Your licence key may have expired or may not have been installed correctly.  
 See Section 3.6.5 in the Essential Introduction for further information.

**7 Accuracy**

The computed singular value decomposition is nearly the exact singular value decomposition for a nearby matrix  $(A + E)$ , where

$$\|E\|_2 = O(\epsilon)\|A\|_2,$$

and  $\epsilon$  is the *machine precision*. In addition, the computed singular vectors are nearly orthogonal to working precision. See Section 4.9 of Anderson *et al.* (1999) for further details.

**8 Parallelism and Performance**

nag\_dgesvd (f08kbc) is threaded by NAG for parallel execution in multithreaded implementations of the NAG Library.

nag\_dgesvd (f08kbc) makes calls to BLAS and/or LAPACK routines, which may be threaded within the vendor library used by this implementation. Consult the documentation for the vendor library for further information.

Please consult the X06 Chapter Introduction for information on how to control and interrogate the OpenMP environment used within this function. Please also consult the Users' Note for your implementation for any additional implementation-specific information.

## 9 Further Comments

The total number of floating-point operations is approximately proportional to  $mn^2$  when  $m > n$  and  $m^2n$  otherwise.

The singular values are returned in descending order.

The complex analogue of this function is nag\_zgesvd (f08kpc).

## 10 Example

This example finds the singular values and left and right singular vectors of the 6 by 4 matrix

$$A = \begin{pmatrix} 2.27 & -1.54 & 1.15 & -1.94 \\ 0.28 & -1.67 & 0.94 & -0.78 \\ -0.48 & -3.09 & 0.99 & -0.21 \\ 1.07 & 1.22 & 0.79 & 0.63 \\ -2.35 & 2.93 & -1.45 & 2.30 \\ 0.62 & -7.39 & 1.03 & -2.57 \end{pmatrix},$$

together with approximate error bounds for the computed singular values and vectors.

The example program for nag\_dgesdd (f08kdc) illustrates finding a singular value decomposition for the case  $m \leq n$ .

### 10.1 Program Text

```

/* nag_dgesvd (f08kbc) Example Program.
 *
 * Copyright 2014 Numerical Algorithms Group.
 *
 * Mark 23, 2011.
 */

#include <math.h>
#include <stdio.h>
#include <nag.h>
#include <nag_stdlib.h>
#include <nagf08.h>
#include <nagf16.h>
#include <nagx02.h>
#include <nagx04.h>

int main(void)
{
    /* Scalars */
    double      alpha, beta, eps, norm, serrbd;
    Integer     exit_status = 0, i, j, m, n, pda, pdd, pdu, pdvt;

    /* Arrays */
    double      *a = 0, *d = 0, *rcondu = 0, *rcondv = 0;
    double      *s = 0, *u = 0, *uerrbd = 0, *verrbd = 0, *vt = 0, *work = 0;

    /* Nag Types */
    NagError    fail;
    Nag_OrderType order;

```

```

#ifdef NAG_COLUMN_MAJOR
#define A(I, J) a[(J - 1) * pda + I - 1]
#define U(I, J) u[(J - 1) * pdu + I - 1]
    order = Nag_ColMajor;
#else
#define A(I, J) a[(I - 1) * pda + J - 1]
#define U(I, J) u[(I - 1) * pdu + J - 1]
    order = Nag_RowMajor;
#endif

    INIT_FAIL(fail);

    printf("nag_dgesvd (f08kbc) Example Program Results\n\n");
    /* Skip heading in data file */
#ifdef _WIN32
    scanf_s("%*[\n]");
#else
    scanf("%*[\n]");
#endif

#ifdef _WIN32
    scanf_s("%"NAG_IFMT%"NAG_IFMT"%*[\n]", &m, &n);
#else
    scanf("%"NAG_IFMT%"NAG_IFMT"%*[\n]", &m, &n);
#endif
    if (m < 0 || n < 0)
    {
        printf("Invalid m or n\n");
        exit_status = 1;
        goto END;
    }

    /* Allocate memory */
    if (!(a      = NAG_ALLOC(m * n, double)) ||
        !(d      = NAG_ALLOC(m * n, double)) ||
        !(rcondu = NAG_ALLOC(n, double)) ||
        !(rcondv = NAG_ALLOC(n, double)) ||
        !(s      = NAG_ALLOC(MIN(m, n), double)) ||
        !(u      = NAG_ALLOC(m * m, double)) ||
        !(uerrbd = NAG_ALLOC(n, double)) ||
        !(verrbd = NAG_ALLOC(n, double)) ||
        !(vt     = NAG_ALLOC(n * n, double)) ||
        !(work   = NAG_ALLOC(MIN(m, n), double)) )
    {
        printf("Allocation failure\n");
        exit_status = -1;
        goto END;
    }

    pdu = m;
    pdvt = n;
#ifdef NAG_COLUMN_MAJOR
    pda = m;
    pdd = m;
#else
    pda = n;
    pdd = n;
#endif

    /* Read the m by n matrix A from data file */
    for (i = 1; i <= m; ++i)
#ifdef _WIN32
        for (j = 1; j <= n; ++j) scanf_s("%lf", &A(i, j));
#else
        for (j = 1; j <= n; ++j) scanf("%lf", &A(i, j));
#endif
#ifdef _WIN32
    scanf_s("%*[\n]");
#else
    scanf("%*[\n]");
#endif

```

```

/* Copy a into d */
for(i = 0; i < m*n; i++) d[i] = a[i];

/* nag_gen_real_mat_print (x04cac)
 * Print real general matrix A.
 */
fflush(stdout);
nag_gen_real_mat_print(order, Nag_GeneralMatrix, Nag_NonUnitDiag, m, n, a,
                       pda, "Matrix A", 0, &fail);
printf("\n");
if (fail.code != NE_NOERROR)
{
    printf("Error from nag_gen_real_mat_print (x04cac).\n%s\n", fail.message);
    exit_status = 1;
    goto END;
}

/* nag_dgesvd (f08kbc).
 * Compute the singular values and left and right singular vectors
 * of A ( $A = U*S*(V^T)$ , m.ge.n)
 */
nag_dgesvd(order, Nag_AllU, Nag_AllVT, m, n, a, pda, s, u, pdu, vt, pdvt,
           work, &fail);
if (fail.code != NE_NOERROR)
{
    printf("Error from nag_dgesvd (f08kbc).\n%s\n", fail.message);
    exit_status = 1;
    goto END;
}

/* U <- U*S */
for(i = 1; i <= m; i++)
    for(j = 1; j <= n; j++) U(i, j) *= s[j-1];

/* nag_dgemm (f16yac):
 * Compute  $D = D - U*S*V^T$  from the factorization of A
 * and store in d */
alpha = -1.0;
beta = 1.0;
nag_dgemm(order, Nag_NoTrans, Nag_NoTrans, m, n, n, alpha, u, pdu, vt, pdvt,
          beta, d, pdd, &fail);
if (fail.code != NE_NOERROR)
{
    printf("Error from nag_dgemm (f16yac).\n%s\n", fail.message);
    exit_status = 1;
    goto END;
}

/* nag_dge_norm (f16rac)
 * Find norm of matrix D and print warning if it is too large.
 */
nag_dge_norm(order, Nag_OneNorm, m, n, d, pdd, &norm, &fail);
if (fail.code != NE_NOERROR)
{
    printf("Error from nag_dge_norm (f16rac).\n%s\n", fail.message);
    exit_status = 1;
    goto END;
}

/* nag_machine_precision (x02ajc): the machine precision. */
eps = nag_machine_precision;
if (norm > pow(eps,0.8))
{
    printf("\nNorm of A-(U*S*V^T) is much greater than 0.\n"
          "Schur factorization has failed.\n");
    exit_status = 1;
    goto END;
}
/* Get the machine precision, eps and compute the approximate
 * error bound for the computed singular values.

```

```

    * Note that for the 2-norm, s[0] = norm(A).
    */
    serrbd = eps * s[0];

    /* Estimate reciprocal condition numbers for the singular vectors using
    * nag_ddisna (f08flc).
    */
    nag_ddisna(Nag_LeftSingVecs, m, n, s, rcondu, &fail);
    nag_ddisna(Nag_RightSingVecs, m, n, s, rcondv, &fail);

    /* Compute the error estimates for the singular vectors */
    for (i = 0; i < n; ++i)
    {
        uerrbd[i] = serrbd / rcondu[i];
        verrbd[i] = serrbd / rcondv[i];
    }

    /* Print the approximate error bounds for the singular values and vectors */
    printf("Error estimate for the singular values\n%11.1e\n", serrbd);

    printf("\nError estimates for the left singular vectors\n");
    for (i = 0; i < n; ++i) printf(" %10.1e%s", uerrbd[i], i%6 == 5?"\n":""");

    printf("\n\nError estimates for the right singular vectors\n");
    for (i = 0; i < n; ++i) printf(" %10.1e%s", verrbd[i], i%6 == 5?"\n":""");
    printf("\n");

END:
    NAG_FREE(a);
    NAG_FREE(d);
    NAG_FREE(rcondu);
    NAG_FREE(rcondv);
    NAG_FREE(s);
    NAG_FREE(u);
    NAG_FREE(uerrbd);
    NAG_FREE(verrbd);
    NAG_FREE(vt);
    NAG_FREE(work);

    return exit_status;
}
#undef A
#undef U

```

## 10.2 Program Data

nag\_dgesvd (f08kbc) Example Program Data

```

    6      4      : m and n

    2.27 -1.54  1.15 -1.94
    0.28 -1.67  0.94 -0.78
   -0.48 -3.09  0.99 -0.21
    1.07  1.22  0.79  0.63
   -2.35  2.93 -1.45  2.30
    0.62 -7.39  1.03 -2.57 : matrix A

```

## 10.3 Program Results

nag\_dgesvd (f08kbc) Example Program Results

```

Matrix A
      1      2      3      4
1     2.2700  -1.5400  1.1500  -1.9400
2     0.2800  -1.6700  0.9400  -0.7800
3    -0.4800  -3.0900  0.9900  -0.2100
4     1.0700   1.2200  0.7900   0.6300
5    -2.3500   2.9300 -1.4500   2.3000
6     0.6200  -7.3900  1.0300  -2.5700

```

Error estimate for the singular values  
1.1e-15

Error estimates for the left singular vectors  
1.8e-16    4.8e-16    1.3e-15    2.2e-15

Error estimates for the right singular vectors  
1.8e-16    4.8e-16    1.3e-15    1.3e-15

---