


Ponderosa Computing
Linear Algebra .NET Class Library



Paul J. McClellan, Ph.D.

7 August 2019

Table of Contents

1	PONDEROSA COMPUTING LINEAR ALGEBRA .NET CLASS LIBRARY	4
2	HOW TO USE THIS CLASS LIBRARY	5
3	NOTATION, IMPLEMENTATION, AND NUMERICAL PRECISION	8
3.1	LINEAR ALGEBRA NOTATION	8
3.2	LIBRARY USES THE CLAPACK IMPLEMENTATION OF LAPACK	9
3.3	LIBRARY COMPUTATIONS USE IEEE 754 DOUBLE PRECISION FORMAT	9
3.3.1	<i>Machine Constants</i>	10
3.3.2	<i>Normalized and Denormalized Numbers</i>	10
3.3.3	<i>Nonfinite Numbers</i>	10
4	LINEARALGEBRA CLASS METHODS.....	11
4.1	MATRIX CREATION AND EXTRACTION	11
4.1.1	<i>Create a Constant Matrix</i>	11
4.1.2	<i>Create a Random Matrix</i>	11
4.1.3	<i>Create an Identity Matrix</i>	12
4.1.4	<i>Create a Diagonal Matrix</i>	12
4.1.5	<i>Extract Diagonal Elements</i>	12
4.1.6	<i>Extract Row Elements</i>	12
4.1.7	<i>Extract Column Elements</i>	12
4.1.8	<i>Extract Rows</i>	13
4.1.9	<i>Extract Columns</i>	13
4.2	MATRIX ADDITION, SUBTRACTION, MULTIPLICATION, AND TRANSPOSE OPERATIONS ..	13
4.2.1	<i>Vector or Matrix Addition</i>	14
4.2.2	<i>Vector or Matrix Subtraction</i>	14
4.2.3	<i>Matrix Multiplication</i>	14
4.2.4	<i>Matrix Transpose Multiplication</i>	15
4.2.5	<i>Matrix Multiply Transpose</i>	15
4.2.6	<i>Transpose</i>	15
4.3	SCALAR-VALUED FUNCTIONS	16
4.3.1	<i>Vector Dot Product</i>	16
4.3.2	<i>1-Norm (Column Norm)</i>	17
4.3.3	<i>Infinity-Norm (Row Norm)</i>	17
4.3.4	<i>Frobenius-Norm</i>	18
4.3.5	<i>2-Norm (Spectral Norm)</i>	18
4.3.6	<i>Rank</i>	19
4.3.7	<i>Spectral Radius</i>	20
4.3.8	<i>Trace</i>	20
4.3.9	<i>1-Norm Inverse Condition Number Estimate</i>	21
4.3.10	<i>Infinity-Norm Inverse Condition Number Estimate</i>	21
4.3.11	<i>Determinant</i>	22
4.4	FULL-RANK SQUARE LINEAR SYSTEM SOLVERS	22

Ponderosa Computing Linear Algebra .NET Class Library

4.4.1	<i>System Solution</i>	22
4.4.2	<i>System Solution with Error Bounds</i>	23
4.5	GENERAL LEAST-SQUARES SYSTEM SOLVERS	25
4.5.1	<i>General Least-Squares System Solvers (via Orthogonal Factorizations)</i>	25
4.5.2	<i>General Least-Squares System Solvers (via SVD)</i>	28
4.6	SINGULAR VALUE DECOMPOSITION	30
4.6.1	<i>Singular Value Decomposition</i>	31
4.6.2	<i>Singular Values</i>	32
4.6.3	<i>Singular Values With Error Bound</i>	32
4.6.4	<i>Left Singular Vectors</i>	33
4.6.5	<i>Right Singular Vectors</i>	34
4.7	EIGENVALUES AND EIGENVECTORS	34
4.7.1	<i>Eigenvalues</i>	35
4.7.2	<i>Eigenvalues with Error Bounds</i>	35
4.7.3	<i>Eigenvalues and Right Eigenvectors</i>	36
4.8	CHOLESKY FACTORIZATION	37
4.8.1	<i>Upper and Lower Cholesky Factorizations</i>	37
5	REFERENCES	39
6	LICENSE NOTICE	40

1 Ponderosa Computing Linear Algebra .NET Class Library

The C# programming language and .NET Framework provide one- and two-dimensional array data types that are well suited to representing linear algebra vectors and matrices. The Ponderosa Computing Linear Algebra .NET class library, **PonderosaComputing.LinearAlgebra.dll**, provides linear algebra vector/matrix metrics and operations using one-dimensional and two-dimensional, double-precision .NET array objects to represent column vectors and matrices, respectively. These metrics and operations are implemented as **LinearAlgebra** class methods using **LAPACK** algorithms.

LAPACK [1,2] is a freely-available, peer-reviewed computational linear algebra software library that provides routines for solving systems of simultaneous linear equations, computing least-squares solutions of linear systems of equations, and computing eigenvalue and singular value decompositions. The associated matrix factorizations (LU, Cholesky, LQ/QR, SVD, Schur, generalized Schur) are also provided, as are related computations such as estimating condition numbers. LAPACK has been used in or as a starting point for implementation of linear algebra computing environments and is a standard by which other libraries and computing environments are often compared.

The Ponderosa Computing Linear Algebra .NET class library provides selected LAPACK linear algebra computations using the public-domain CLAPACK library from the Netlib Repository [3,4]. The CLAPACK library is a machine-translation of the LAPACK Fortran library to C code. The **LinearAlgebra** class methods are implemented using C++ Interop in C++/CLI to wrap the CLAPACK functions so they can be accessed by code that is authored in C# or another .NET Framework language [5].

This document describes version 1.2 of the Ponderosa Computing Linear Algebra .NET class library. It was built using the Microsoft .NET Framework 4.5.2.

2 How to Use this Class Library

We illustrate here how to use the Ponderosa Computing Linear Algebra .NET class library in a C# project using Microsoft Visual Studio 2012.

The Ponderosa Computing Linear Algebra .NET class library installer installs by default the class library DLL, **PonderosaComputing.LinearAlgebra.dll**, in the folder C:\Program Files (x86)\Ponderosa Computing\LinearAlgebra.NET. If needed, the installer will also attempt to install the .NET Framework 4.5.2 from the web.

To use the class methods of this class library first add a reference to the library in your Visual Studio C# project:

1. In Visual Studio click the menu item “Project” and “Add Reference...”.
2. Click the “Browse...” button on the lower right and locate the file **PonderosaComputing.LinearAlgebra.dll**. By default the installer installs this file in C:\Program Files (x86)\Ponderosa Computing\LinearAlgebra.NET.
3. Click on the file name and click the “Add” button on the lower right of the file dialog.
4. Click the “OK” button.

Then create a **PonderosaComputing.LinearAlgebra** class instance and call the desired class method through this instance. **LinearAlgebraExceptions** can also be caught.

Here is example C# code using this library:

```
using System;
using PonderosaComputing;

namespace PcLinAlgCLTest
{
    class Program
    {
        static void WriteVector(double[] Vec)
        {
            for (int i = 0; i < Vec.GetLength(0); ++i)
            {
                Console.Write("{0}  ", Vec[i]);
            }
            Console.WriteLine();
            Console.WriteLine();
        }

        static void WriteMatrix(double[,] Mat)
        {
            for (int i = 0; i < Mat.GetLength(0); ++i)
            {
                for (int j = 0; j < Mat.GetLength(1); ++j)
                {
                    Console.Write("{0}  ", Mat[i, j]);
                }
                Console.WriteLine();
            }
        }
    }
}
```

Ponderosa Computing Linear Algebra .NET Class Library

```
    }
    Console.WriteLine();
}

static void Main(string[] args)
{
    LinearAlgebra la = new LinearAlgebra();

    // Exception thrown
    double[,] MS = {{1,2,4},{1,2,4},{1,2,4}};
    double[] v = {1,2,3};
    try
    {
        Console.WriteLine("*** Test error handling ***");
        Console.WriteLine("");

        Console.WriteLine("Singular matrix M");
        WriteMatrix(MS);
        Console.WriteLine("rank = {0}", la.Rank(MS));
        Console.WriteLine("determinant = {0}", la.Determinant(MS));
        Console.WriteLine();

        Console.WriteLine("vector v");
        WriteVector(v);

        Console.WriteLine("Matrix-vector product u = M * v");
        double[] u = la.Multiply(MS, v);
        WriteVector(u);

        Console.WriteLine("Try to solve M * x = u for x");
        Console.WriteLine("Expect exception");
        Console.WriteLine();
        double[] x = la.SolveFullRankLinearSystem(MS, u);
        Console.WriteLine("Expected not to reach here!");
        WriteVector(x);
    }
    catch (LinearAlgebraException e)
    {
        Console.WriteLine("Exception caught: {0}", e.Message);
    }
    finally
    {
        Console.WriteLine();
    }

    // metrics
    double[,] M33 = {{1,2,3},{1,4,9},{1,8,27}};
    double[,] M43 = {{1,2,3},{1,4,9},{1,8,27},{1,16,81}};
    try
    {
        Console.WriteLine("*** Solve full rank square system ***");
        Console.WriteLine("");

        Console.WriteLine("Full rank matrix M:");
        WriteMatrix(M33);
        Console.WriteLine("rank = {0}", la.Rank(M33));
        Console.WriteLine("determinant = {0}", la.Determinant(M33));
    }
}
```

Ponderosa Computing Linear Algebra .NET Class Library

```
        Console.WriteLine();

        Console.WriteLine("vector v");
        WriteVector(v);

        Console.WriteLine("Matrix-vector product u = M * v");
        double[] u = la.Multiply(M33, v);
        WriteVector(u);

        Console.WriteLine("Solve M * x = u for x");
        double[] x = la.SolveFullRankLinearSystem(M33, u);
        WriteVector(x);

        Console.WriteLine("*** Solve overdetermined system ***");
        Console.WriteLine("");

        Console.WriteLine("matrix H:");
        WriteMatrix(M43);
        Console.WriteLine("rank = {0}", la.Rank(M43));
        Console.WriteLine("");

        Console.WriteLine("Matrix-vector product w = H * v");
        double[] w = la.Multiply(M43, v);
        WriteVector(w);

        Console.WriteLine("Solve H * y = w for y with error bound:");
        double[] ye = la.SolveLsqLinearSystemEB(M43, w);
        WriteVector(ye);
    }
    catch (LinearAlgebraException e)
    {
        Console.WriteLine("Exception caught: {0}", e.Message);
    }
    finally
    {
        Console.WriteLine();
    }

    Console.ReadLine();
}
}
```

3 Notation, Implementation, and Numerical Precision

3.1 Linear Algebra Notation

A linear algebra **column vector** is a one-dimensional array of elements consisting of numbers, symbols, or expressions arranged in a column. A column vector of m elements consists of **elements** x_i , where i is the row location of the element:

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{bmatrix} .$$

The Ponderosa Computing Linear Algebra .NET class library uses double-precision one-dimensional .NET arrays to represent column vectors. A column vector can also be represented by a double-precision two-dimensional .NET array object having one column.

A linear algebra **row vector** is a one-dimensional array of elements consisting of numbers, symbols, or expressions arranged in a row. A row vector of n elements consists of **elements** y_i , where i is the column location of the element:

$$\mathbf{y} = [y_1 \quad y_2 \quad \cdots \quad y_n] .$$

This class library uses double-precision two-dimensional .NET arrays with one row to represent row vectors.

A linear algebra **matrix** is a two-dimensional rectangular array of elements consisting of numbers, symbols, or expressions arranged in rows and columns. A matrix of m rows and n columns consists of **elements** x_{ij} , where i is the row location and j is the column location of the element:

$$\mathbf{X} = \begin{bmatrix} x_{11} & x_{12} & \cdots & x_{1n} \\ x_{21} & x_{22} & \cdots & x_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ x_{m1} & x_{m2} & \cdots & x_{mn} \end{bmatrix} .$$

This class library uses double-precision two-dimensional .NET arrays to represent matrices.

The **transpose** operation $\mathbf{Y} = \mathbf{X}^T$ on an $m \times n$ matrix \mathbf{X} creates an $n \times m$ matrix \mathbf{Y} with rows and columns interchanged:

$$\mathbf{Y} = \mathbf{X}^T = \begin{bmatrix} y_{11} & y_{12} & \cdots & y_{1m} \\ y_{21} & y_{22} & \cdots & y_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ y_{n1} & y_{n2} & \cdots & y_{nm} \end{bmatrix} , \text{ where } y_{ij} = x_{ji} .$$

The transpose operation $\mathbf{y} = \mathbf{x}^T$ on an m-element column vector \mathbf{x} creates an m-element row vector:

$$\mathbf{x}^T = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{bmatrix}^T = [x_1 \quad x_2 \quad \dots \quad x_m].$$

This class library will return the transpose of a column vector as a row vector represented by a double-precision two-dimensional .NET array with one row.

Numerical linear algebra operations are defined for matrices and vectors with all elements defined on the domain of complex numbers \mathbb{C} . The Ponderosa Computing Linear Algebra .NET class library provides linear algebra vector/matrix metrics and operations using double-precision one-dimensional and two-dimensional .NET array objects to represent vector and matrix arguments and return values. Since these array objects are declared as double-precision the elements of these vectors and matrices are defined on the domain of real numbers $\mathbb{R} = (-\infty, +\infty)$.

3.2 Library Uses the CLAPACK Implementation of LAPACK

LAPACK [1] is a freely-available, peer-reviewed numerical linear algebra software package that provides routines for solving systems of simultaneous linear equations, computing least-squares solutions of linear systems of equations, and computing eigenvalue and singular value decompositions. The associated matrix factorizations (LU, Cholesky, LQ/QR, SVD, Schur, generalized Schur) are also provided, as are related computations such as estimating condition numbers.

The Netlib Repository provides a cross-platform Fortran source distribution of LAPACK [2] and a C source distributions of CLAPACK [3]. Instructions and tools for building CLAPACK on the Windows platform are available from the University of Tennessee Innovative Computing Laboratory [4]. The Ponderosa Computing Linear Algebra .NET class library uses the CLAPACK implementation of the LAPACK software package.

3.3 Library Computations use IEEE 754 Double Precision Format

The numerical linear algebra computations are all defined for matrix and vector elements defined on the domain of real numbers $\mathbb{R} = (-\infty, +\infty)$. To describe limiting behavior of linear algebra computations it can be useful to define them for matrix and vector elements defined on the domain of the **affinely extended real numbers** $[-\infty, +\infty]$, which adds the elements $+\infty$ (positive infinity) and $-\infty$ (negative infinity) to the real numbers.

The Ponderosa Computing Linear Algebra .NET class library implements its functions using the double precision format defined by the IEEE Standard 754 for Binary Floating-Point Arithmetic [6]. This format includes representations of signed zeros and normalized, denormalized, and nonfinite (infinite and indeterminate) floating point numbers.

3.3.1 Machine Constants

We define here some double precision constants that appear later in this document.

eps : The relative machine precision, the distance from 1.0 to the next largest double-precision number. This number is $eps = \text{dlamch}(P) = 2^{-52} = 2.2204460492503131e - 016$.

safemin : The minimum positive floating point value such that $1/safemin$ does not overflow. $safemin = \text{dlamch}(S) = 2.2250738585072014e - 308$.

3.3.2 Normalized and Denormalized Numbers

The Ponderosa Computing Linear Algebra .NET class library supports normalized IEEE 754 double precision numbers with 53 binary digits of precision (equivalent to nearly 16 decimal digits of precision).

The library also supports IEEE 754 denormalized numbers. As an example, for the matrix

$$A = \begin{bmatrix} 1E - 160 & 0 \\ 0 & 1E - 160 \end{bmatrix},$$

the class method `LinearAlgebra.Determinant()` returns the denormalized value $9.99988867182683E-321$.

3.3.3 Nonfinite Numbers

The CLAPACK implementation does not reliably handle nonfinite (infinite and indeterminate) floating point matrix or vector elements. For example, some computational loops are skipped when a factor is zero under the assumption that the skipped loop would have no impact on the computed results [7]. But this may fail to propagate nonfinite values or fail to detect invalid operations.

Version 1.2 of the Ponderosa Computing Linear Algebra .NET class library does not provide special handling of nonfinite floating point matrix or vector elements, and its class methods might not provide reliable results for method arguments containing such elements.

4 LinearAlgebra Class Methods

The Ponderosa Computing Linear Algebra .NET class library implements its functions using version 3.2.1 of CLAPACK [3]. This implementation of CLAPACK was machine translated by Netlib from Fortran 77 to ANSI C using the f2c tool and version 3.2.1 of LAPACK [2]. The Ponderosa Computing Linear Algebra .NET class library uses the reference BLAS library included with this CLAPACK distribution.

4.1 Matrix Creation and Extraction

The Ponderosa Computing Linear Algebra .NET class library provides the following class methods for creating matrices and extracting the diagonals of a matrix:

Create a constant matrix	GenerateConstant()
Create a random matrix	GenerateRandom()
Create an identity matrix	GenerateIdentity()
Create a diagonal matrix	GenerateDiagonal()
Extract the diagonal elements of a matrix	ExtractDiagonals()
Extract a row vector from a matrix	ExtractRow()
Extract a column vector from a matrix	ExtractColumn()
Extract rows from a matrix	ExtractRows()
Extract columns from a matrix	ExtractColumns()

4.1.1 Create a Constant Matrix

The class method

LinearAlgebra.GenerateConstant(m, n, value)

returns an $m \times n$ matrix with entries all equal to value. If $m < 1$ or $n < 1$ then this method throws an **eBadParamError** LinearAlgebraException.

4.1.2 Create a Random Matrix

The class method

LinearAlgebra.GenerateRandom(m, n, minimum, maximum)

returns an $m \times n$ matrix with random entries in [minimum, maximum). If $m < 1$ or $n < 1$ or maximum < minimum then this method throws an **eBadParamError** LinearAlgebraException.

This class method generates uniformly distributed random elements in the interval [minimum, maximum] using a Mersenne Twister pseudo-random generator of 32-bit numbers with a state size of 19937 bits provided by the C++11 standard library. This generator is initialized when the **PonderosaComputing.LinearAlgebra** class instance is constructed and a class method is first invoked through this instance.

4.1.3 Create an Identity Matrix

The class method

LinearAlgebra.GenerateIdentity(n)

returns an $n \times n$ identity matrix, I_n , of order n . If $n < 1$ then this method throws an **eBadParamError** LinearAlgebraException.

4.1.4 Create a Diagonal Matrix

The class method

LinearAlgebra.GenerateDiagonal(d)

returns an $n \times n$ diagonal matrix D where n is the size of the argument vector d containing the diagonal elements of the returned diagonal matrix.

4.1.5 Extract Diagonal Elements

The class method

LinearAlgebra.ExtractDiagonals(A)

returns the $\min(m, n)$ main diagonal elements of an $m \times n$ matrix A as a $\min(m, n)$ -element vector.

4.1.6 Extract Row Elements

The class method

LinearAlgebra.ExtractRow(A, r)

returns the n elements of the r -th row of an $m \times n$ matrix A as a n -element column vector v . The row number r is base 0. If $r < 0$ or $r \geq m$ then this method throws an **eBadParamError** LinearAlgebraException.

4.1.7 Extract Column Elements

The class method

LinearAlgebra.ExtractColumn (A, c)

returns the m elements of the c -th column of an $m \times n$ matrix A as an m -element vector v . The column number c is base 0. If $c < 0$ or $c \geq n$ then this method throws an **eBadParamError** LinearAlgebraException.

4.1.8 Extract Rows

The class method

LinearAlgebra.ExtractRows (A, r₁, r₂)

returns the rows r_1 through r_2 of an $m \times n$ matrix A as a $(r_2 - r_1 + 1) \times n$ matrix. The row numbers r_1 and r_2 are base 0. If $r_1 < 0$ or $r_1 > r_2$ or $r_2 \geq m$ then this method throws an **eBadParamError** LinearAlgebraException.

4.1.9 Extract Columns

The class method

LinearAlgebra.ExtractColumns (A, c₁, c₂)

returns the columns c_1 through c_2 of an $m \times n$ matrix A as a $m \times (c_2 - c_1 + 1)$ matrix. The column numbers c_1 and c_2 are base 0. If $c_1 < 0$ or $c_1 > c_2$ or $c_2 \geq m$ then this method throws an **eBadParamError** LinearAlgebraException.

4.2 Matrix Addition, Subtraction, Multiplication, and Transpose Operations

The Ponderosa Computing Linear Algebra .NET class library provides the following class methods for vector and matrix addition and subtraction and for matrix multiplication and transpose operations:

Vector or Matrix addition	Add ()
Vector or Matrix subtraction	Subtract ()
Matrix multiplication	Multiply ()
Matrix transpose multiplication	TransposeMultiply ()
Matrix multiply transpose	MultiplyTranspose ()
Transpose	Transpose ()

4.2.1 Vector or Matrix Addition

The class method

LinearAlgebra.Add(a,b)

returns the element-wise sum of n-element column vectors **a** and **b** as an n-element column vector **c** = **a** + **b**. If the number of elements of **a** does not match the number of elements of **b** then this method throws an **eBadParamError** LinearAlgebraException.

The class method

LinearAlgebra.Add(A,B)

returns the element-wise sum of m x n matrices **A** and **B** as an m x n matrix **C** = **A** + **B**. If the number of rows and columns of **A** do not match the number of rows and columns of **B** then this method throws an **eBadParamError** LinearAlgebraException.

4.2.2 Vector or Matrix Subtraction

The class method

LinearAlgebra.Subtract(a,b)

returns the element-wise difference of n-element column vectors **a** and **b** as an n-element column vector **c** = **a** - **b**. If the number of elements of **a** does not match the number of elements of **b** then this method throws an **eBadParamError** LinearAlgebraException.

The class method

LinearAlgebra.Subtract(A,B)

returns the element-wise difference of m x n matrices **A** and **B** as an m x n matrix **C** = **A** - **B**. If the number of rows and columns of **A** do not match the number of rows and columns of **B** then this method throws an **eBadParamError** LinearAlgebraException.

4.2.3 Matrix Multiplication

The class method

LinearAlgebra.Multiply(A,b)

returns the product of an m x n matrix **A** and a n-element column vector **b** as an m-element column vector **v** = **A** **b**. If the number of columns of **A** does not match the number of elements of **b** then this method throws an **eBadParamError** LinearAlgebraException.

The class method

LinearAlgebra.Multiply(A,B)

returns the product of an $m \times n$ matrix A and a $n \times p$ matrix B as an $m \times p$ matrix $C = AB$. If the number of columns of A does not match the number of rows of B then this method throws an **eBadParamError** LinearAlgebraException.

These methods compute the matrix product using routine `dgemm_()` from the CLAPACK v 3.2.1 BLAS package.

4.2.4 Matrix Transpose Multiplication

The class method

LinearAlgebra.TransposeMultiply(A,b)

returns the product of the transpose of an $n \times m$ matrix A and a n -element column vector b as an m -element column vector $v = A^T b$. If the number of rows of A does not match the number of elements of b then this method throws an **eBadParamError** LinearAlgebraException.

The class method

LinearAlgebra.TransposeMultiply(A,B)

returns the product of an $m \times n$ matrix A and a $n \times p$ matrix B as an $m \times p$ matrix $C = A^T B$. If the number of columns of A does not match the number of rows of B then this method throws an **eBadParamError** LinearAlgebraException.

These methods use routine `dgemm_()` from the CLAPACK v 3.2.1 BLAS package.

4.2.5 Matrix Multiply Transpose

The class method

LinearAlgebra.MultiplyTranspose(A,B)

returns the product of an $m \times n$ matrix A and the transpose of a $p \times n$ matrix B as an $m \times p$ matrix $C = AB^T$. If the number of columns of A does not match the number of columns of B then this method throws an **eBadParamError** LinearAlgebraException.

This method uses routine `dgemm_()` from the CLAPACK v 3.2.1 BLAS package.

4.2.6 Transpose

The class method

LinearAlgebra.Transpose(v)

returns the transpose of an m-element column vector v as the 1 x m matrix representing the row vector $u = v^T$.

The class method

LinearAlgebra.Transpose(A)

returns the transpose of an m x n matrix A as the n x m matrix A^T .

4.3 Scalar-Valued Functions

The Ponderosa Computing Linear Algebra .NET class library provides the following scalar-valued vector and matrix methods:

Vector dot product	Dot()
1-norm (column norm) of a vector or matrix	OneNorm()
Infinity-norm (row norm) of a vector or matrix	InfinityNorm()
Frobenius norm of a vector or matrix	FrobeniusNorm()
2-norm (spectral norm) of a vector or matrix	TwoNorm()
Rank of a matrix	Rank()
Spectral radius of a symmetric matrix	SpectralRadius()
Trace of a square matrix	Trace()
1-norm inverse condition number estimate of a square matrix	InverseOneNormConditionNumberEstimate()
Infinity-norm inverse condition number estimate of a square matrix	InverseInfinityNormConditionNumberEstimate()
Determinant of a square matrix	Determinant()

4.3.1 Vector Dot Product

The class method

LinearAlgebra.Dot(*u*,*v*)

returns the dot product of two *n*-element column vectors ***u*** and ***v***. This is the sum of the products of the corresponding elements of ***u*** and ***v***:

$$\mathbf{u} \cdot \mathbf{v} = \sum_{i=1}^n u_i * v_i$$

These methods compute the dot product using routine `ddot_()` from the CLAPACK v 3.2.1 BLAS package.

4.3.2 1-Norm (Column Norm)

The class methods

LinearAlgebra.OneNorm(*A*)

LinearAlgebra.OneNorm(*v*)

return the 1-norm (column norm, $\|\mathbf{A}\|_1$, $\|\mathbf{v}\|_1$) of a matrix ***A*** or column vector ***v***, respectively.

This is the maximum absolute column sum of the elements of an *m* x *n* matrix ***A***:

$$\|\mathbf{A}\|_1 = \max_{1 \leq j \leq n} \sum_{i=1}^m |a_{ij}|$$

It is the absolute column sum of the elements of an *n* element column vector ***v***:

$$\|\mathbf{v}\|_1 = \sum_{i=1}^n |v_i|$$

This method computes the 1-norm using the routine `dlange_()` from CLAPACK v 3.2.1.

4.3.3 Infinity-Norm (Row Norm)

The class methods

LinearAlgebra.InfinityNorm(*A*)

LinearAlgebra.InfinityNorm(*v*)

return the infinity-norm (row norm, $\|\mathbf{A}\|_\infty$, $\|\mathbf{v}\|_\infty$) of a matrix ***A*** or column vector ***v***, respectively.

This is the maximum absolute row sum of the elements of an $m \times n$ matrix \mathbf{A} :

$$\|\mathbf{A}\|_{\infty} = \max_{1 \leq i \leq m} \sum_{j=1}^n |a_{ij}|$$

It is the maximum absolute value of the elements of an n element column vector \mathbf{v} :

$$\|\mathbf{v}\|_{\infty} = \max_{1 \leq i \leq n} |v_i|$$

This method computes the infinity-norm using routine `dlange_()` from CLAPACK v 3.2.1.

4.3.4 Frobenius-Norm

The class methods

`LinearAlgebra.FrobeniusNorm(A)`

`LinearAlgebra.FrobeniusNorm(v)`

return the Frobenius-norm ($\|\mathbf{A}\|_F$, $\|\mathbf{v}\|_F$) of a matrix \mathbf{A} or vector \mathbf{v} , respectively.

This is the square root of the sum of absolute squares (root mean square) of the elements of an $m \times n$ matrix \mathbf{A} :

$$\|\mathbf{A}\|_F = \sqrt{\sum_{i=1}^m \sum_{j=1}^n |a_{ij}|^2}$$

It is the square root of the sum of absolute squares (root mean square) of the elements of an n vector \mathbf{v} :

$$\|\mathbf{v}\|_F = \sqrt{\sum_{i=1}^n |v_i|^2}$$

This method computes the Frobenius norm using routine `dlange_()` from CLAPACK v 3.2.1.

4.3.5 2-Norm (Spectral Norm)

The class methods

`LinearAlgebra.TwoNorm(A)`

LinearAlgebra.TwoNorm(v)

return the 2-norm (spectral norm, $\|A\|_2$, $\|v\|_2$) of a matrix **A** or column vector **v**, respectively.

This is the largest singular value of an m x n matrix **A**:

$$\|A\|_2 = \sigma_{\max}(A) .$$

It is the square root of the sum of absolute squares (root mean square) of the elements of an n vector **v**:

$$\|v\|_F = \sqrt{\sum_{i=1}^n |v_i|^2}$$

This method computes the min(m, n) singular values of a m x n matrix **A** using routine `dgesvd_()` from CLAPACK v 3.2.1.

Routine `dgesvd_()` uses routine `dbdsqr_()` to compute the singular values of an (upper or lower) bidiagonal matrix using the implicit zero-shift QR algorithm. If this algorithm fails to find all the singular values of **A** then this method throws an **InternalError** LinearAlgebraException.

This method computes the 2-norm of a vector using routine `dlange_()` from CLAPACK v 3.2.1.

4.3.6 Rank

The **column rank** of an m x n matrix **A** is the maximum number of linearly independent column vectors of **A**. The **row rank** of an m x n matrix **A** is the maximum number of linearly independent row vectors of **A**. The column rank and the row rank are equal, and this is called the **rank** of the matrix **A**. The rank of an m x n matrix **A** is also the number of nonzero singular values of **A**.

The class method

LinearAlgebra.Rank(A)

returns the rank of a m x n matrix **A**. This method computes the min(m, n) singular values of **A** using routine `dgesvd_()` from CLAPACK v 3.2.1. The rank is then determined as the number of computed singular values that are significantly greater than zero.

The Ponderosa Computing Linear Algebra .NET class library uses the threshold value

$$threshold = \max(sfmin, \max(m, n) * eps * \sigma_{\max}(A)) .$$

A computed singular value is considered significantly greater than zero if it exceeds this threshold. Here *safemin* is the minimum positive floating point value such that $1/safemin$ does not overflow, *eps* is the relative machine precision, and $\sigma_{\max}(\mathbf{A})$ is the largest singular value of \mathbf{A} .

Routine `dgesvd_()` uses routine `dbdsqr_()` to compute the singular values of an (upper or lower) bidiagonal matrix using the implicit zero-shift QR algorithm. If this algorithm fails to find all the singular values of \mathbf{A} then this method throws an **eInternalError** `LinearAlgebraException`.

4.3.7 Spectral Radius

Let $\lambda_1, \lambda_2, \dots, \lambda_n$ be the (real) eigenvalues of an $n \times n$ symmetric (real) matrix \mathbf{A} . Then the spectral radius of \mathbf{A} is:

$$\rho(\mathbf{A}) = \max_i(|\lambda_i|).$$

The class method

LinearAlgebra.SpectralRadius(A)

returns the spectral radius of a symmetric matrix \mathbf{A} . If \mathbf{A} is not symmetric then this method throws an **eBadParamError** `LinearAlgebraException`.

This method computes the eigenvalues of \mathbf{A} using routine `dsyev_()` from CLAPACK v 3.2.1. Routine `dsyev_()` uses routine `dsterf_()` which computes all eigenvalues of a symmetric tridiagonal matrix using the Pal-Walker-Kahan variant of the QL or QR algorithm. If this algorithm fails to find all of the eigenvalues of \mathbf{A} in at most $30*n$ iterations then this method throws an **eInternalError** `LinearAlgebraException`.

4.3.8 Trace

The trace of an $n \times n$ matrix \mathbf{A} is the sum of the main diagonal elements of \mathbf{A} :

$$tr(\mathbf{A}) = \sum_{i=1}^n a_{ii}$$

The class method

LinearAlgebra.Trace(A)

returns the trace of a square matrix \mathbf{A} directly from its definition. If \mathbf{A} is not square then this method throws an **eBadParamError** `LinearAlgebraException`.

4.3.9 1-Norm Inverse Condition Number Estimate

The class method

LinearAlgebra.InverseOneNormConditionNumberEstimate (A)

returns the inverse of a 1-norm condition number estimate of a square matrix A . If A is not square this method throws an **eBadParamError** LinearAlgebraException.

This method starts by computing $\|A\|_1$ using routine `dlange_()` from CLAPACK v 3.2.1 and returns zero if $\|A\|_1 = 0$.

This method then uses routine `dgetrf_()` from CLAPACK v 3.2.1 to compute the LU factorization $A = PLU$ using partial pivoting with row interchanges. If routine `dgetrf_()` determines the matrix A is exactly singular then this method returns 0.

Otherwise, this method estimates $\|A^{-1}\|_1$ using the LU factorization and routine `dgecon_()` from CLAPACK v 3.2.1 and returns zero if the $\|A^{-1}\|_1$ estimate is zero. Otherwise this method returns the inverse condition number estimate

$$c = \frac{1}{\|A\|_1 \|A^{-1}\|_1}$$

LinearAlgebra.InverseOneNormConditionNumberEstimate () returns the inverse of the 1-norm condition number estimate of a matrix, rather than the condition number estimate of the matrix, itself, to provide a zero return value for singular matrices.

4.3.10 Infinity-Norm Inverse Condition Number Estimate

The class method

LinearAlgebra.InverseInfinityNormConditionNumberEstimate (A)

returns the inverse of an infinity-norm condition number estimate of a square matrix A . If A is not square this method throws an **eBadParamError** LinearAlgebraException.

This method starts by computing $\|A\|_\infty$ using routine `dlange_()` from CLAPACK v 3.2.1 and returns zero if $\|A\|_\infty = 0$.

This method then uses routine `dgetrf_()` from CLAPACK v 3.2.1 to compute the LU factorization $A = PLU$ using partial pivoting with row interchanges. If routine `dgetrf_()` determines the matrix A is exactly singular then this method returns 0.

Otherwise, This method estimates $\|A^{-1}\|_{\infty}$ using the LU factorization and routine `dgecon_()` from CLAPACK v 3.2.1 and returns zero if the $\|A^{-1}\|_{\infty}$ estimate is zero. Otherwise This method returns the inverse condition number estimate

$$c = \frac{1}{\|A\|_{\infty}\|A^{-1}\|_{\infty}}$$

LinearAlgebra.InverseInfinityNormConditionNumberEstimate() returns the inverse of the infinity-norm condition number estimate of a matrix, rather than the condition number estimate of the matrix, itself, to provide a zero return value for singular matrices.

4.3.11 Determinant

The class method

LinearAlgebra.Determinant(A)

returns the determinant of a square matrix **A**. If **A** is not square then this method throws an **eBadParamError** LinearAlgebraException.

This method computes the determinant of **A** by using routine `dgetrf_()` from CLAPACK v 3.2.1 to compute the LU factorization $A = PLU$ using partial pivoting with row interchanges. If routine `dgetrf_()` determines the matrix **A** is exactly singular then this method returns 0. Otherwise, This method accumulates the product of the diagonal entries of **U** with sign adjustment according to pivot row interchanges and scaling to avoid intermediate overflow.

4.4 Full-Rank Square Linear System Solvers

The Ponderosa Computing Linear Algebra .NET class library provides the following full-rank square linear system solver methods:

Full rank square system solvers

`SolveFullRankLinearSystem()`
`SolveFullRankLinearSystemEB()`

4.4.1 System Solution

The class method

LinearAlgebra.SolveFullRankLinearSystem(A,b)

returns the solution vector **x** to a full-rank real linear system:

$$Ax = b$$

using equilibration and iterative refinement as needed. Here \mathbf{A} is an $n \times n$ matrix and \mathbf{b} is a n -element column vector. The solution vector \mathbf{x} is a n -element column vector. If \mathbf{A} is not square or if vector \mathbf{b} does not have n elements this method throws an **eBadParamError** LinearAlgebraException.

The class method

LinearAlgebra.SolveFullRankLinearSystem(A,B)

returns the solution matrix \mathbf{X} to a full-rank real linear system:

$$\mathbf{A} \mathbf{X} = \mathbf{B}$$

using equilibration and iterative refinement as needed. Here \mathbf{A} is an $n \times n$ matrix and \mathbf{B} is a $n \times p$ matrix. The solution matrix \mathbf{X} is $n \times p$. If \mathbf{A} is not square or if matrix \mathbf{B} does not have n rows this method throws an **eBadParamError** LinearAlgebraException.

These methods compute the solution matrix \mathbf{X} by using routine `dgesvx_()` from CLAPACK v 3.2.1, using equilibration and iterative refinement as needed. If routine `dgesvx_()` determines the matrix \mathbf{A} is exactly singular then these methods throw an **eInternalError** LinearAlgebraException.

4.4.2 System Solution with Error Bounds

The class method

LinearAlgebra.SolveFullRankLinearSystemEB(A,b)

returns the solution vector \mathbf{x} and forward error bound estimate e to a full-rank real linear system:

$$\mathbf{A} \mathbf{x} = \mathbf{b}$$

using equilibration and iterative refinement as needed. Here \mathbf{A} is an $n \times n$ matrix and \mathbf{b} is a n -element column vector. The solution vector \mathbf{x} is a n -element column vector and the forward error bound estimate e is a scalar. If \mathbf{A} is not square or if vector \mathbf{b} does not have n elements this method throws an **eBadParamError** LinearAlgebraException.

The computed solution vector \mathbf{x} is returned as the first n elements of a $(n+1)$ -element augmented solution vector \mathbf{x}_A . The forward error bound estimate e is returned as the last element of \mathbf{x}_A .

$$\mathbf{x}_A = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \\ e \end{bmatrix}.$$

The forward error bound estimate e bounds the relative error in the computed solution. If \mathbf{x} is the computed solution and \mathbf{x}_t is the true solution, then the relative error is bounded by:

$$\frac{\|\mathbf{x} - \mathbf{x}_t\|_\infty}{\|\mathbf{x}\|_\infty} \leq e$$

That is, e is an estimated upper bound for the magnitude of the largest element in $\mathbf{x} - \mathbf{x}_t$ divided by the magnitude of the largest element in \mathbf{x} . This estimate is almost always a slight overestimate of the true error.

The class method

LinearAlgebra.SolveFullRankLinearSystemEB(A,B)

returns the solution matrix \mathbf{X} and forward error bound estimates e to a full-rank real linear system:

$$\mathbf{A}\mathbf{X} = \mathbf{B}$$

using equilibration and iterative refinement as needed. Here \mathbf{A} is an $n \times n$ matrix and \mathbf{B} is a $n \times p$ matrix. The solution matrix \mathbf{X} is $n \times p$ and the forward error bound estimates e is a row vector. If \mathbf{A} is not square or if matrix \mathbf{B} does not have n rows this method throws an **eBadParamError** LinearAlgebraException.

The computed matrix solution \mathbf{X} is returned as the first n rows of a $(n+1) \times p$ augmented solution matrix \mathbf{X}_A . The forward error bound estimate row vector e is returned as the last row of \mathbf{X}_A .

$$\mathbf{X}_A = \begin{bmatrix} x_{11} & x_{12} & \cdots & x_{1p} \\ x_{21} & x_{22} & \cdots & x_{2p} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \cdots & x_{np} \\ e_1 & e_2 & \cdots & e_p \end{bmatrix}.$$

For each column j of the augmented solution matrix, the forward error bound estimate e_j bounds the relative error in the computed solution column \mathbf{x}_j . If \mathbf{x}_j is the computed solution column and \mathbf{x}_{jt} is the true solution column, then the relative error for that column is bounded by:

$$\frac{\|\mathbf{x}_j - \mathbf{x}_{jt}\|_\infty}{\|\mathbf{x}_j\|_\infty} \leq e_j$$

That is, e_j is an estimated upper bound for the magnitude of the largest element in $\mathbf{x}_j - \mathbf{x}_{jt}$ divided by the magnitude of the largest element in \mathbf{x}_j . This estimate is almost always a slight overestimate of the true error.

These methods compute the augmented solution matrix X by using routine `dgesvx_()` from CLAPACK v 3.2.1, using equilibration and iterative refinement as needed. If routine `dgesvx_()` determines the matrix A is exactly singular then these methods throw an `eInternalError` `LinearAlgebraException`.

4.5 General Least-Squares System Solvers

The Ponderosa Computing Linear Algebra .NET class library provides the following real linear least squares system solver methods:

General least-squares system solvers (via LQ/QR)	<code>SolveLsqLinearSystem()</code> <code>SolveLsqLinearSystemEB()</code>
General least-squares system solvers (via SVD)	<code>SolveLsqLinearSystemSvd()</code> <code>SolveLsqLinearSystemSvdEB()</code>

4.5.1 General Least-Squares System Solvers (via Orthogonal Factorizations)

4.5.1.1 System Solution

The class method

`LinearAlgebra.SolveLsqLinearSystem(A, b)`

returns the minimum-norm solution vector x to a real linear least squares system:

$$\min_x \|b - Ax\|_2, \text{ with minimum } \|x\|_2.$$

Here A is an $m \times n$ matrix which may be rank-deficient and b is an m -element column vector. The solution is the minimum-2-norm n -element column vector x achieving the minimum 2-norm residual. This method uses a complete orthogonal factorization of A . If the number of rows of A does not match the number of elements of b then this method throws an `eBadParamError` `LinearAlgebraException`.

Multiple real linear least squares systems sharing the same system matrix A can be solved in a single method call. The p right hand side m -element column vectors b_j for the least squares systems:

$$\min_x \|b_j - Ax_j\|_2, \text{ with minimum } \|x_j\|_2, \quad j = 1 \dots p$$

are stored as the columns of a $m \times p$ matrix B .

The class method:

LinearAlgebra.SolveLsqLinearSystem(A,B)

returns the minimum-norm solution vector \mathbf{x}_j for the j-th real linear least squares system as the j-th column of the $n \times p$ solution matrix \mathbf{X} . If the number of rows of \mathbf{A} does not match the number of rows of \mathbf{B} then this method throws an **eBadParamError** LinearAlgebraException.

These methods compute the solution vector \mathbf{x} and matrix \mathbf{X} by using routine `dgelsy_()` from CLAPACK v 3.2.1. This routine first computes a QR factorization of \mathbf{A} with column pivoting. It then determines the effective rank r of \mathbf{A} and then refactors a reduced system into an $r \times r$ triangular system which is then solved. If routine `dgelsy_()` determines the matrix \mathbf{A} has effective rank 0 then these methods return the zero solution vector \mathbf{x} or matrix \mathbf{X} .

4.5.1.2 System Solution with Error Bounds

The class method

LinearAlgebra.SolveLsqLinearSystemEB(A,b)

returns the minimum-norm solution vector \mathbf{x} and forward error bound estimate e to a real linear least squares system:

$$\min_{\mathbf{x}} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2, \text{ with minimum } \|\mathbf{x}\|_2.$$

Here \mathbf{A} is an $m \times n$ matrix which may be rank-deficient and \mathbf{b} is an m -element column vector. The solution is the minimum-2-norm n -element column vector \mathbf{x} achieving the minimum 2-norm residual. The forward error bound estimate e is a scalar. This method uses a complete orthogonal factorization of \mathbf{A} . If the number of rows of \mathbf{A} does not match the number of elements of \mathbf{b} then this method throws an **eBadParamError** LinearAlgebraException.

The computed solution vector \mathbf{x} is returned as the first n elements of a $(n+1)$ -element augmented solution vector \mathbf{x}_A . If the matrix \mathbf{A} has $m \geq n$ (we have an overdetermined system) with full column rank then this method also returns a forward error bound estimate e as the last element of \mathbf{x}_A . Otherwise this method returns the value -1 as the last element of \mathbf{x}_A .

$$\mathbf{x}_A = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \\ e \end{bmatrix}.$$

The forward error bound estimate e bounds the relative error in the computed solution. If \mathbf{x} is the computed solution and \mathbf{x}_t is the true solution, then the relative error is bounded by:

$$\frac{\|\mathbf{x} - \mathbf{x}_t\|_2}{\|\mathbf{x}_t\|_2} \leq e$$

That is, e is an estimated upper bound for the root mean square of $\mathbf{x} - \mathbf{x}_t$ divided by the root mean square of \mathbf{x}_t . This estimate is almost always a slight overestimate of the true error.

Multiple real linear least squares systems sharing the same system matrix \mathbf{A} can be solved with forward error bound estimates in a single method call. The p right hand side m -element column vectors \mathbf{b}_j for the least squares systems:

$$\min_x \|\mathbf{b}_j - \mathbf{A}\mathbf{x}_j\|_2, \text{ with minimum } \|\mathbf{x}_j\|_2, \quad j = 1 \dots p$$

are stored as the columns of a $m \times p$ matrix \mathbf{B} .

The class method:

LinearAlgebra.SolveLsqLinearSystemEB(A,B)

returns the minimum-norm solution vector \mathbf{x}_j for the j -th real linear least squares system as the j -th column of the first n rows of the augmented $(n+1) \times p$ augmented solution matrix \mathbf{X}_A . If the matrix \mathbf{A} has $m \geq n$ (we have an overdetermined system) with full column rank then this method also returns a forward error bound estimate e_j returned as the last element of column j of \mathbf{X}_A . Otherwise this method returns the value -1 as the last element of column j of \mathbf{X}_A .

$$\mathbf{X}_A = \begin{bmatrix} x_{11} & x_{12} & \cdots & x_{1p} \\ x_{21} & x_{22} & \cdots & x_{2p} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \cdots & x_{np} \\ e_1 & e_2 & \cdots & e_p \end{bmatrix}.$$

If the number of rows of \mathbf{A} does not match the number of rows of \mathbf{B} then this method throws an **eBadParamError** LinearAlgebraException.

For each computed solution vector \mathbf{x}_j , the forward error bound estimate e_j bounds the relative error in the solution \mathbf{x}_j . If \mathbf{x}_j is the computed solution vector and \mathbf{x}_{jt} is the true solution vector, then the relative error for that solution vector is bounded by:

$$\frac{\|\mathbf{x}_j - \mathbf{x}_{jt}\|_2}{\|\mathbf{x}_{jt}\|_2} \leq e$$

That is, e_j is an estimated upper bound for the root mean square of $\mathbf{x}_j - \mathbf{x}_{jt}$ divided by the root mean square of \mathbf{x}_{jt} . This estimate is almost always a slight overestimate of the true error.

These methods compute the solution vector \mathbf{x} and matrix \mathbf{X} by using routine `dgelsy_()` from CLAPACK v 3.2.1. This routine first computes a QR factorization of \mathbf{A} with column pivoting. It then determines the effective rank r of \mathbf{A} and then refactors a reduced system into an $r \times r$

triangular system which is then solved. If routine `dgelsy_()` determines the matrix \mathbf{A} has effective rank 0 then these methods return the zero solution vector \mathbf{x} or matrix \mathbf{X} .

The error bounds for the full column rank overdetermined case are computed in the manner described here: <http://www.netlib.org/lapack/lug/node82.html>.

4.5.2 General Least-Squares System Solvers (via SVD)

4.5.2.1 System Solution

The class method

LinearAlgebra.SolveSvdLinearSystem(A, b)

returns the minimum-norm solution vector \mathbf{x} to a real linear least squares system:

$$\min_{\mathbf{x}} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2, \text{ with minimum } \|\mathbf{x}\|_2.$$

Here \mathbf{A} is an $m \times n$ matrix which may be rank-deficient and \mathbf{b} is an m -element column vector. The solution is the minimum-2-norm n -element column vector \mathbf{x} achieving the minimum 2-norm residual. This method uses the singular value decomposition of \mathbf{A} . If the number of rows of \mathbf{A} does not match the number of elements of \mathbf{b} then this method throws an **eBadParamError** `LinearAlgebraException`.

Multiple real linear least squares systems sharing the same system matrix \mathbf{A} can be solved in a single method call. The p right hand side m -element column vectors \mathbf{b}_j for the least squares systems:

$$\min_{\mathbf{x}} \|\mathbf{b}_j - \mathbf{A}\mathbf{x}_j\|_2, \text{ with minimum } \|\mathbf{x}_j\|_2, \quad j = 1 \dots p$$

are stored as the columns of a $m \times p$ matrix \mathbf{B} .

The class method:

LinearAlgebra.SolveLsqLinearSystem(A, B)

returns the minimum-norm solution vector \mathbf{x}_j for the j -th real linear least squares system as the j -th column of the $n \times p$ solution matrix \mathbf{X} . If the number of rows of \mathbf{A} does not match the number of rows of \mathbf{B} then this method throws an **eBadParamError** `LinearAlgebraException`.

These methods compute the solution vector \mathbf{x} and matrix \mathbf{X} by using routine `dgelss_()` from CLAPACK v 3.2.1. This routine first computes the singular value decomposition of \mathbf{A} . It then determines the effective rank r of \mathbf{A} , zeros out the insignificant portion of the SVD, and solves the reduced system. If routine `dgelss_()` determines the matrix \mathbf{A} has effective rank 0 then these methods return the zero solution vector \mathbf{x} or matrix \mathbf{X} .

4.5.2.2 System Solution with Error Bounds

The class method

LinearAlgebra.SolveSvdLinearSystemEB(A,b)

returns the minimum-norm solution vector \mathbf{x} and forward error bound estimate e to a real linear least squares system:

$$\min_{\mathbf{x}} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2, \text{ with minimum } \|\mathbf{x}\|_2.$$

Here \mathbf{A} is an $m \times n$ matrix which may be rank-deficient and \mathbf{b} is an m -element column vector. The solution is the minimum-2-norm n -element column vector \mathbf{x} achieving the minimum 2-norm residual. The forward error bound estimate e is a scalar. This method uses the singular value decomposition of \mathbf{A} . If the number of rows of \mathbf{A} does not match the number of elements of \mathbf{b} then this method throws an **eBadParamError** LinearAlgebraException.

The computed solution vector \mathbf{x} is returned as the first n elements of a $(n+1)$ -element augmented solution vector \mathbf{x}_A . If the matrix \mathbf{A} has $m \geq n$ (we have an overdetermined system) with full column rank then this method also returns a forward error bound estimate e as the last element of \mathbf{x}_A . Otherwise this method returns the value -1 as the last element of \mathbf{x}_A .

$$\mathbf{x}_A = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \\ e \end{bmatrix}.$$

The forward error bound estimate e bounds the relative error in the computed solution. If \mathbf{x} is the computed solution and \mathbf{x}_t is the true solution, then the relative error is bounded by:

$$\frac{\|\mathbf{x} - \mathbf{x}_t\|_2}{\|\mathbf{x}_t\|_2} \leq e$$

That is, e is an estimated upper bound for the root mean square of $\mathbf{x} - \mathbf{x}_t$ divided by the root mean square of \mathbf{x}_t . This estimate is almost always a slight overestimate of the true error.

Multiple real linear least squares systems sharing the same system matrix \mathbf{A} can be solved with forward error bound estimates in a single method call. The p right hand side m -element column vectors \mathbf{b}_j for the least squares systems:

$$\min_{\mathbf{x}} \|\mathbf{b}_j - \mathbf{A}\mathbf{x}_j\|_2, \text{ with minimum } \|\mathbf{x}_j\|_2, \quad j = 1 \dots p$$

are stored as the columns of a $m \times p$ matrix \mathbf{B} .

The class method:

LinearAlgebra.SolveSvdLinearSystemEB(A,B)

returns the minimum-norm solution vector \mathbf{x}_j for the j -th real linear least squares system as the j -th column of the first n rows of the augmented $(n+1) \times p$ augmented solution matrix \mathbf{X}_A . If the matrix \mathbf{A} has $m \geq n$ (we have an overdetermined system) with full column rank then this method also returns a forward error bound estimate e_j returned as the last element of column j of \mathbf{X}_A . Otherwise this method returns the value -1 as the last element of column j of \mathbf{X}_A .

$$\mathbf{X}_A = \begin{bmatrix} x_{11} & x_{12} & \cdots & x_{1p} \\ x_{21} & x_{22} & \cdots & x_{2p} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \cdots & x_{np} \\ e_1 & e_2 & \cdots & e_p \end{bmatrix}.$$

If the number of rows of \mathbf{A} does not match the number of rows of \mathbf{B} then this method throws an **eBadParamError** LinearAlgebraException.

For each computed solution vector \mathbf{x}_j , the forward error bound estimate e_j bounds the relative error in the solution \mathbf{x}_j . If \mathbf{x}_j is the computed solution vector and \mathbf{x}_{jt} is the true solution vector, then the relative error for that solution vector is bounded by:

$$\frac{\|\mathbf{x}_j - \mathbf{x}_{jt}\|_2}{\|\mathbf{x}_{jt}\|_2} \leq e$$

That is, e_j is an estimated upper bound for the root mean square of $\mathbf{x}_j - \mathbf{x}_{jt}$ divided by the root mean square of \mathbf{x}_{jt} . This estimate is almost always a slight overestimate of the true error.

These methods compute the solution vector \mathbf{x} and matrix \mathbf{X} by using routine `dgelss_()` from CLAPACK v 3.2.1. This routine first computes the singular value decomposition of \mathbf{A} . It then determines the effective rank r of \mathbf{A} , zeros out the insignificant portion of the SVD, and solves the reduced system. If routine `dgelss_()` determines the matrix \mathbf{A} has effective rank 0 then these methods return the zero solution vector \mathbf{x} or matrix \mathbf{X} .

The error bounds for the full column rank overdetermined case are computed in the manner described here: <http://www.netlib.org/lapack/lug/node82.html>.

4.6 Singular Value Decomposition

The **singular value decomposition** of an $m \times n$ (real) matrix \mathbf{A} is a factorization of \mathbf{A} into the form:

$$\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$$

where \mathbf{U} is a $m \times m$ matrix, \mathbf{V}^T is an $n \times n$ matrix, and $\mathbf{\Sigma}$ is a $m \times n$ rectangular diagonal matrix with $p = \min(m,n)$ nonnegative entries on the main diagonal.

The nonnegative diagonal entries of $\mathbf{\Sigma}$ are the **singular values** of \mathbf{A} , $\{\sigma_1, \sigma_2, \dots, \sigma_{\min(m,n)}\}$, arranged in descending order: $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{\min(m,n)} \geq 0$.

The matrices \mathbf{U} and \mathbf{V} are orthonormal:

$$\mathbf{U}^T \mathbf{U} = \mathbf{U} \mathbf{U}^T = \mathbf{I}_m, \quad \mathbf{V}^T \mathbf{V} = \mathbf{V} \mathbf{V}^T = \mathbf{I}_n,$$

where \mathbf{I}_m and \mathbf{I}_n denote the identity matrices of order m and n , respectively. The m columns of \mathbf{U} , $\{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m\}$, are called the **left singular vectors** of \mathbf{A} . The n columns of \mathbf{V} , $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$, are called the **right singular vectors** of \mathbf{A} .

Since $\mathbf{\Sigma}$ is a $m \times n$ rectangular diagonal matrix with $p = \min(m,n)$, the singular value decomposition can be represented as the **thin singular value decomposition**:

$$\mathbf{A} = \mathbf{U}_p \mathbf{\Sigma}_p \mathbf{V}_p^T,$$

Where \mathbf{U}_p is the $m \times p$ matrix consisting of the first p columns of \mathbf{U} , $\mathbf{\Sigma}_p$ is the $p \times p$ upper diagonal portion of $\mathbf{\Sigma}$, and \mathbf{V}_p is the $n \times p$ matrix consisting of the first p columns of \mathbf{V} . The p columns of \mathbf{U}_p and the p columns of \mathbf{V}_p are orthonormal:

$$\mathbf{U}_p^T \mathbf{U}_p = \mathbf{I}_p, \quad \mathbf{V}_p^T \mathbf{V}_p = \mathbf{I}_p,$$

where \mathbf{I}_p denote the identity matrix of order p .

The Ponderosa Computing Linear Algebra .NET class library provides the following methods computing the thin singular value decomposition or parts thereof of a matrix:

Singular value decomposition	SingularValueDecomposition ()
Singular values	SingularValues () SingularValuesEB ()
Left or right singular vectors	SingularValuesLeftVectors () SingularValuesRightVectors ()

4.6.1 Singular Value Decomposition

The class method

LinearAlgebra.SingularValueDecomposition (A)

returns the thin singular value decomposition of an $m \times n$ matrix \mathbf{A} . This function returns the $p = \min(m,n)$ singular values, their associated left singular vectors, and their associated right singular vectors of \mathbf{A} in the p columns of a $(m+n+1) \times p$ matrix \mathbf{S} :

$$\mathbf{S} = \begin{bmatrix} \sigma_1 & \sigma_2 & \cdots & \sigma_p \\ u_{11} & u_{21} & \cdots & u_{p1} \\ \vdots & \vdots & \ddots & \vdots \\ u_{1m} & u_{2m} & \cdots & u_{pm} \\ v_{11} & v_{21} & \cdots & v_{p1} \\ \vdots & \vdots & \ddots & \vdots \\ v_{1n} & v_{2n} & \cdots & v_{pn} \end{bmatrix}$$

The first row of \mathbf{S} contains the p singular values of \mathbf{A} , $\{\sigma_1, \sigma_2, \dots, \sigma_p\}$, in descending order. The columns of the next m rows contain the first p left singular vectors, $\{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_p\}$. The columns of the last n rows contain the first p right singular vectors, $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p\}$.

This method computes the singular value decomposition of \mathbf{A} using routine `dgesvd_()` from CLAPACK v 3.2.1. Routine `dgesvd_()` uses routine `dbdsqr_()` to compute the singular values and singular vectors of an (upper or lower) bidiagonal matrix using the implicit zero-shift QR algorithm. If this algorithm fails to find all the singular values of \mathbf{A} then this method throws an **eInternalError** `LinearAlgebraException`.

4.6.2 Singular Values

The class method

LinearAlgebra.SingularValues(A)

returns the $p = \min(m,n)$ singular values of a $m \times n$ matrix \mathbf{A} , $\{\sigma_1, \sigma_2, \dots, \sigma_p\}$, in descending order as an p -element column vector \mathbf{s} :

$$\mathbf{s} = \begin{bmatrix} \sigma_1 \\ \sigma_2 \\ \vdots \\ \sigma_p \end{bmatrix}.$$

This method computes the singular values of \mathbf{A} using routine `dgesvd_()` from CLAPACK v 3.2.1. Routine `dgesvd_()` uses routine `dbdsqr_()` to compute the singular values of an (upper or lower) bidiagonal matrix using the implicit zero-shift QR algorithm. If this algorithm fails to find all the singular values of \mathbf{A} then this method throws an **eInternalError** `LinearAlgebraException`.

4.6.3 Singular Values With Error Bound

The class method

LinearAlgebra.SingularValuesEB(A)

returns the $p = \min(m,n)$ singular values of a $m \times n$ matrix \mathbf{A} , $\{\sigma_1, \sigma_2, \dots, \sigma_p\}$, in descending order followed by a forward error bound estimate e as an $(p+1)$ -element column vector \mathbf{s}_A :

$$\mathbf{s}_A = \begin{bmatrix} \sigma_1 \\ \sigma_2 \\ \vdots \\ \sigma_p \\ e \end{bmatrix} .$$

The forward error bound estimate e bounds the absolute error in the computed singular values. If σ_i is the i -th computed singular value and σ_{it} is the i -th true singular value, then the absolute error in σ_i is bounded by:

$$|\sigma_i - \sigma_{it}| \leq e$$

This estimate is almost always a slight overestimate of the true error.

This method computes the singular values of \mathbf{A} using routine `dgesvd_()` from CLAPACK v 3.2.1. Routine `dgesvd_()` uses routine `dbdsqr_()` to compute the singular values of an (upper or lower) bidiagonal matrix using the implicit zero-shift QR algorithm. If this algorithm fails to find all the singular values of \mathbf{A} then this method throws an **eInternalError** `LinearAlgebraException`.

The singular values forward error bound e is:

$$e = \max(sfmin, \max(m,n) * eps * \sigma_1) .$$

This error bound is described here: <http://www.netlib.org/lapack/lug/node97.html>, but this methods uses $p(m,n) = \max(m,n)$ in place of $p(m,n) = 1$ as the "modestly growing function of n " and uses machine epsilon `dlamch_("P")` in place of `dlamch_("E")`. The computation of e is similar to that found in the example program: <http://www.nag.com/lapack-ex/node128.html> .

4.6.4 Left Singular Vectors

The class method

LinearAlgebra.SingularValuesLeftVectors(A)

returns the $p = \min(m,n)$ singular values and first p left singular vectors of a $m \times n$ matrix \mathbf{A} in the p columns of a $(m+1) \times p$ matrix \mathbf{S} :

$$\mathbf{S} = \begin{bmatrix} \sigma_1 & \sigma_2 & \cdots & \sigma_p \\ u_{11} & u_{21} & \cdots & u_{p1} \\ \vdots & \vdots & \ddots & \vdots \\ u_{1m} & u_{2m} & \cdots & u_{pm} \end{bmatrix}$$

The first row of \mathbf{S} contains the p singular values of \mathbf{A} , $\{\sigma_1, \sigma_2, \dots, \sigma_p\}$, in descending order. The columns of the remaining m rows contain the first p left singular vectors, $\{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_p\}$.

This method computes the singular value decomposition of \mathbf{A} using routine `dgesvd_()` from CLAPACK v 3.2.1. Routine `dgesvd_()` uses routine `dbdsqr_()` to compute the singular values and singular vectors of an (upper or lower) bidiagonal matrix using the implicit zero-shift QR algorithm. If this algorithm fails to find all the singular values of \mathbf{A} then this method throws an **eInternalError** `LinearAlgebraException`.

4.6.5 Right Singular Vectors

The class method

`LinearAlgebra.SingularValuesRightVectors(A)`

returns the $p = \min(m,n)$ singular values and first p right singular vectors of a $m \times n$ matrix \mathbf{A} in the p rows of a $p \times (n+1)$ matrix \mathbf{S} :

$$\mathbf{S} = \begin{bmatrix} \sigma_1 & v_{11} & v_{12} & \cdots & v_{1n} \\ \sigma_2 & v_{21} & v_{22} & \cdots & v_{2n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \sigma_p & v_{p1} & v_{p2} & \cdots & v_{pn} \end{bmatrix}$$

The first column of \mathbf{S} contains the p singular values of \mathbf{A} in descending order. The rows of the remaining n columns of \mathbf{S} contain the first p right singular vectors, $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p\}$.

This method computes the singular value decomposition of \mathbf{A} using routine `dgesvd_()` from CLAPACK v 3.2.1. Routine `dgesvd_()` uses routine `dbdsqr_()` to compute the singular values and singular vectors of an (upper or lower) bidiagonal matrix using the implicit zero-shift QR algorithm. If this algorithm fails to find all the singular values of \mathbf{A} then this method throws an **eInternalError** `LinearAlgebraException`.

4.7 Eigenvalues and Eigenvectors

A **right eigenvector** of an $n \times n$ symmetric (real) matrix \mathbf{A} is a nonzero n -vector \mathbf{v} such that the matrix product:

$$\mathbf{A}\mathbf{v} = \lambda\mathbf{v} .$$

Here λ is a real number (scalar) called the **eigenvalue** of \mathbf{A} corresponding to the right eigenvector \mathbf{v} . The set of all eigenvectors of \mathbf{A} , $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$, each paired with its corresponding eigenvalue, $\{\lambda_1, \lambda_2, \dots, \lambda_n\}$, is called the **eigensystem** of the matrix \mathbf{A} . It can be expressed as:

$$\mathbf{A} \mathbf{V} = \mathbf{V} \mathbf{\Lambda}.$$

Here $\mathbf{\Lambda}$ is a $n \times n$ diagonal matrix with diagonal entries $\{\lambda_1, \lambda_2, \dots, \lambda_n\}$ and \mathbf{V} is a $n \times n$ matrix with columns $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$.

The Ponderosa Computing Linear Algebra .NET class library provides the following methods computing the eigenvalues and right eigenvectors of a symmetric (real) matrix:

Eigenvalues	Eigenvalues() EigenvaluesEB()
Eigenvalues and right eigenvectors	EigenvaluesRightEigenvectors()

4.7.1 Eigenvalues

The class method

LinearAlgebra.Eigenvalues(A)

returns the eigenvalues of an $n \times n$ symmetric (real) matrix \mathbf{A} , $\{\lambda_1, \lambda_2, \dots, \lambda_n\}$, in ascending order, $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$, as an n -element column vector $\boldsymbol{\lambda}$:

$$\boldsymbol{\lambda} = \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \vdots \\ \lambda_n \end{bmatrix}.$$

If \mathbf{A} is not symmetric then this method throws an **eBadParamError** LinearAlgebraException.

This method computes the eigenvalues of \mathbf{A} using routine `dsyev_()` from CLAPACK v 3.2.1. Routine `dsyev_()` uses routine `dsterf_()` which computes all eigenvalues of a symmetric tridiagonal matrix using the Pal-Walker-Kahan variant of the QL or QR algorithm. If this algorithm fails to find all of the eigenvalues of \mathbf{A} in at most $30*n$ iterations then this method throws an **eInternalError** LinearAlgebraException.

4.7.2 Eigenvalues with Error Bounds

The class method

LinearAlgebra.EigenvaluesEB(A)

returns the eigenvalues of an $n \times n$ symmetric (real) matrix \mathbf{A} , $\{\lambda_1, \lambda_2, \dots, \lambda_n\}$, in ascending order, $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$, followed by a forward error bound estimate e as an $(n+1)$ -element column vector $\boldsymbol{\lambda}_A$:

$$\boldsymbol{\lambda}_A = \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \vdots \\ \lambda_n \\ e \end{bmatrix} .$$

If \mathbf{A} is not symmetric then this method throws an **eBadParamError** LinearAlgebraException.

The forward error bound estimate e bounds the absolute error in the computed eigenvalues. If λ_i is the i -th computed eigenvalue and λ_{it} is the i -th true eigenvalue, then the absolute error in λ_i is bounded by:

$$|\lambda_i - \lambda_{it}| \leq e$$

This estimate is almost always a slight overestimate of the true error.

This method computes the eigenvalues of \mathbf{A} using routine `dsyev_()` from CLAPACK v 3.2.1. Routine `dsyev_()` uses routine `dsterf_()` which computes all eigenvalues of a symmetric tridiagonal matrix using the Pal-Walker-Kahan variant of the QL or QR algorithm. If this algorithm fails to find all of the eigenvalues of \mathbf{A} in at most $30*n$ iterations then this method throws an **eInternalError** LinearAlgebraException.

The eigenvalues forward error bound e is:

$$e = \max(\text{sfmin} , n * \text{eps} * \max_i(|\lambda_i|)) .$$

This error bound is described here: <http://www.netlib.org/lapack/lug/node90.html>, but this method uses $p(n) = n$ in place of $p(n) = 1$ as the "modestly growing function of n " and uses machine epsilon `dlamch_("P")` in place of `dlamch_("E")`. The computation of e is similar to that found in the example program: <http://www.nag.com/lapack-ex/node71.html>.

4.7.3 Eigenvalues and Right Eigenvectors

The class method

LinearAlgebra.EigenvaluesRightEigenvectors (A)

returns the eigenvalues and right eigenvectors of an $n \times n$ symmetric (real) matrix \mathbf{A} in the n columns of the $(n+1) \times n$ matrix \mathbf{V}_A .

$$\mathbf{V}_A = \begin{bmatrix} \lambda_1 & \lambda_2 & \cdots & \lambda_n \\ v_{11} & v_{21} & \cdots & v_{n1} \\ \vdots & \vdots & \ddots & \vdots \\ v_{1n} & v_{2n} & \cdots & v_{nn} \end{bmatrix}$$

The first row of \mathbf{V}_A contains the n eigenvalues of \mathbf{A} , $\{\lambda_1, \lambda_2, \dots, \lambda_n\}$, in ascending order, $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$. The columns of the remaining n rows of \mathbf{V}_A contain the corresponding n right eigenvectors of \mathbf{A} , $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$.

If \mathbf{A} is not symmetric then this method throws an **eBadParamError** LinearAlgebraException.

This method computes the eigenvalues of \mathbf{A} using routine `dsyev_()` from CLAPACK v 3.2.1. Routine `dsyev_()` uses routine `dsterf_()` which computes all eigenvalues of a symmetric tridiagonal matrix using the Pal-Walker-Kahan variant of the QL or QR algorithm. If this algorithm fails to find all of the eigenvalues of \mathbf{A} in at most $30*n$ iterations then this method throws an **eInternalError** LinearAlgebraException.

4.8 Cholesky Factorization

The **upper triangular Cholesky factorization** (or decomposition) of an $n \times n$ symmetric positive definite matrix \mathbf{A} is the unique factorization of \mathbf{A} into

$$\mathbf{A} = \mathbf{U}^T \mathbf{U}$$

where \mathbf{U} is an $n \times n$ upper triangular matrix with positive entries on the diagonal.

The **lower triangular Cholesky factorization** (or decomposition) of an $n \times n$ symmetric positive definite matrix \mathbf{A} is the unique factorization of \mathbf{A} into

$$\mathbf{A} = \mathbf{L} \mathbf{L}^T$$

where \mathbf{L} is an $n \times n$ lower triangular matrix with positive entries on the diagonal.

The Ponderosa Computing Linear Algebra .NET class library provides the following methods computing the Cholesky factorizations of a symmetric positive definite matrix:

Upper Cholesky factorization	<code>UpperCholeskyFactorization()</code>
Lower Cholesky factorization	<code>LowerCholeskyFactorization()</code>

4.8.1 Upper and Lower Cholesky Factorizations

The class method

`LinearAlgebra.UpperCholeskyFactorization(A)`

returns the upper triangular Cholesky factorization of an $n \times n$ symmetric positive definite matrix A as an $n \times n$ upper triangular matrix U with positive entries on the diagonal.

The class method

LinearAlgebra.LowerCholeskyFactorization(A)

returns the lower triangular Cholesky factorization of an $n \times n$ symmetric positive definite matrix A as an $n \times n$ lower triangular matrix L with positive entries on the diagonal.

These methods compute a Cholesky factorization of an $n \times n$ symmetric positive definite matrix A using routine `dpotrf_()` from CLAPACK v 3.2.1. If this algorithm determines the the matrix A is not positive definite then these methods throw an **eBadParamError** LinearAlgebraException.

5 References

- [1] Anderson, E. et al., LAPACK Users' Guide, Third Edition (Society for Industrial and Applied Mathematics, Philadelphia, PA, 1999) ISBN 0-89871-447-8.
- [2] LAPACK on the Netlib Repository at UTK and ORNL. <http://www.netlib.org/lapack/>
- [3] CLAPACK (f2c'ed version of LAPACK) version 3.2.1 on the Netlib Repository at UTK and ORNL. <http://www.netlib.org/clapack/>. See clapack-3.2.1-CMAKE.tgz.
- [4] CLAPACK for Windows at the University of Tennessee Innovative Computing Laboratory. <http://icl.cs.utk.edu/lapack-for-windows/clapack/>
- [5] How to: Wrap Native Class for Use by C#. <https://msdn.microsoft.com/en-us/library/ms235281.aspx>
- [6] [Wikipedia: IEEE floating point](#). This discussion includes references to the standard publications.
- [7] LAPACK/ScaLAPACK Development Forum. <http://icl.cs.utk.edu/lapack-forum/>

6 License Notice

The Ponderosa Computing Linear Algebra .NET class library was built using the CLAPACK implementation of LAPACK, a [freely-available software package](http://www.netlib.org/lapack) from netlib at <http://www.netlib.org/lapack>. The license used for the LAPACK software is [the modified BSD license](#):

Copyright (c) 1992-2013 The University of Tennessee and The University of Tennessee Research Foundation. All rights reserved.

Copyright (c) 2000-2013 The University of California Berkeley. All rights reserved.

Copyright (c) 2006-2013 The University of Colorado Denver. All rights reserved.

\$COPYRIGHT\$

Additional copyrights may follow

\$HEADER\$

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer listed in this license in the documentation and/or other materials provided with the distribution.
- Neither the name of the copyright holders nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

The copyright holders provide no reassurances that the source code provided does not infringe any patent, copyright, or any other intellectual property rights of third parties. The copyright holders disclaim any liability to any recipient for claims brought against recipient by any third party for infringement of that parties intellectual property rights.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.