
PyArmadillo: an alternative approach to linear algebra in Python

Jason Rumengan ^{†‡}, Terry Yue Zhuo [◇], Conrad Sanderson ^{†*}

[†] Data61/CSIRO, Australia

[‡] Queensland University of Technology, Australia

[◇] University of New South Wales, Australia

^{*} Griffith University, Australia



PyArmadillo is a streamlined linear algebra library for the Python language, with an emphasis on ease of use. It aims to provide a high-level syntax and functionality deliberately similar to Matlab/Octave [8], allowing mathematical operations to be expressed in a familiar and natural manner. PyArmadillo provides objects for matrices and cubes, as well as over 200 associated functions for manipulating data stored in the objects. Integer, floating point and complex numbers are supported. Various matrix factorisations are provided through integration with LAPACK [2], or one of its high performance drop-in replacements such as Intel MKL [6] or OpenBLAS [12].

While frameworks such as NumPy [5] and SciPy [11] are available for Python, they tend to be overly verbose and somewhat cumbersome to use in comparison to Matlab/Octave. From a linear algebra point of view, the issues include: (i) a focus on operations involving multi-dimensional arrays rather than matrices, (ii) nested organisation of functions which increases code verbosity (which in turn can lead to reduced user productivity), and (iii) syntax that considerably differs from Matlab. PyArmadillo aims to address these issues by primarily focusing on linear algebra, having all functions accessible in one flat structure, and providing an API that closely follows Matlab.

PyArmadillo relies on Armadillo [9, 10] for the underlying C++ implementation of matrix objects and associated functions, as well as on pybind11 [7] for interfacing C++ and Python. Due to its expressiveness and relatively straightforward use, pybind11 was selected over other interfacing approaches such as Boost.Python [1] and manually writing C++ extensions for Python. Armadillo was selected due to its API having close similarity to Matlab, as well as previous success with RcppArmadillo, a widely used bridge between the R language and C++ [4]. PyArmadillo is open source software available under the permissive Apache License 2.0 [3], making it useful in both open-source and proprietary (closed-source) contexts.

A short PyArmadillo based Python program is shown in Figure 1. An overview of the available functionality (as of PyArmadillo version 0.500) is given in Tables 1 to 8. Table 1 briefly describes the member functions and variables of the main matrix class; Table 2 lists the main subset of overloaded Python operators; Table 3 outlines matrix decompositions and equation solvers; Table 4 overviews functions for generating vectors with various sequences; Table 5 lists the main forms of general functions of matrices; Table 6 lists various element-wise functions of matrices; Table 7 summarises the set of functions and classes focused on statistics; Table 8 lists functions specific to signal and image processing. Lastly, Table 9 provides examples of Matlab/Octave syntax and conceptually corresponding PyArmadillo syntax.

PyArmadillo can be obtained from:

- <https://pyarma.sourceforge.io>

If you use PyArmadillo in your research and/or software, we would appreciate a citation to this document. Citations are useful for the continued development and maintenance of the library. Please cite as:

- Jason Rumengan, Terry Yue Zhuo, Conrad Sanderson.
PyArmadillo: an alternative approach to linear algebra in Python.
Technical Report, Data61/CSIRO, 2021.
-

```

from pyarma import *

A = mat(4, 5, fill.ones)
B = mat(4, 5, fill.randu)

C = A * B.t()

C.print("C:")

```

Figure 1: A short PyArmadillo based Python program.

Table 1: Subset of member functions and variables of the *mat* class, the main matrix object in PyArmadillo.

Function/Variable	Description
.n_rows	number of rows (read only)
.n_cols	number of columns (read only)
.n_elem	total number of elements (read only)
[i]	access the <i>i</i> -th element, assuming a column-by-column layout
[r, c]	access the element at row <i>r</i> and column <i>c</i>
.in_range(i)	test whether the <i>i</i> -th element can be accessed
.in_range(r, c)	test whether the element at row <i>r</i> and column <i>c</i> can be accessed
.reset()	set the number of elements to zero
.copy_size(A)	set the size to be the same as matrix <i>A</i>
.set_size(n_rows, n_cols)	change size to specified dimensions, without preserving data (fast)
.reshape(n_rows, n_cols)	change size to specified dimensions, with elements copied column-wise (slow)
.resize(n_rows, n_cols)	change size to specified dimensions, while preserving elements & their layout (slow)
.ones(n_rows, n_cols)	set all elements to one, optionally first resizing to specified dimensions
.zeros(n_rows, n_cols)	as above, but set all elements to zero
.randu(n_rows, n_cols)	as above, but set elements to uniformly distributed random values in [0,1] interval
.randn(n_rows, n_cols)	as above, but use a Gaussian/normal distribution with $\mu = 0$ and $\sigma = 1$
.fill(k)	set all elements to be equal to <i>k</i>
.for_each(lambda val : ...)	for each element, pass its value to a lambda function
.is_empty()	test whether there are no elements
.is_finite()	test whether all elements are finite
.is_square()	test whether the matrix is square
.is_vec()	test whether the matrix is a vector
.is_sorted()	test whether the matrix is sorted
.has_inf()	test whether any element is $\pm\infty$
.has_nan()	test whether any element is not-a-number
iter(A)	iterator of matrix <i>A</i>
iterator(A, i, j)	iterator of matrix <i>A</i> from the <i>i</i> -th element to the <i>j</i> -th
col_iter(A, i, j)	iterator of matrix <i>A</i> from the <i>i</i> -th column to the <i>j</i> -th
row_iter(A, i, j)	iterator of matrix <i>A</i> from the <i>i</i> -th row to the <i>j</i> -th
.print(header)	print elements, with an optional text header
.save(name, format)	store matrix in the specified file, optionally specifying storage format
.load(name, format)	retrieve matrix from the specified file, optionally specifying format
[diag, k]	read/write access to <i>k</i> -th diagonal
[i, :]	read/write access to row <i>i</i>
[:, i]	read/write access to column <i>i</i>
[a:b, :]	read/write access to submatrix, spanning from row <i>a</i> to row <i>b</i>
[:, c:d]	read/write access to submatrix, spanning from column <i>c</i> to column <i>d</i>
[a:b, c:d]	read/write access to submatrix spanning rows <i>a</i> to <i>b</i> and columns <i>c</i> to <i>d</i>
[p, q, size(A)]	read/write access to submatrix starting at row <i>p</i> and col <i>q</i> with size same as matrix <i>A</i>
[vector_of_row_indices, :]	read/write access to rows corresponding to the specified indices
[:, vector_of_col_indices]	read/write access to columns corresponding to the specified indices
[vector_of_indices]	read/write access to matrix elements corresponding to the specified indices
.swap_rows(p, q)	swap the contents of specified rows
.swap_cols(p, q)	swap the contents of specified columns
.insert_rows(row, X)	insert a copy of <i>X</i> at the specified row
.insert_cols(col, X)	insert a copy of <i>X</i> at the specified column
.shed_rows(first_row, last_row)	remove the specified range of rows
.shed_cols(first_col, last_col)	remove the specified range of columns
.min()	return minimum value
.max()	return maximum value
.index_min()	return index of minimum value
.index_max()	return index of maximum value

Table 2: Subset of matrix operations involving overloaded Python operators.

Operation	Description
$A - k$	subtract scalar k from all elements in matrix A
$k - A$	subtract each element in matrix A from scalar k
$A + k, k + A$	add scalar k to all elements in matrix A
$A * k, k * A$	multiply matrix A by scalar k
$A + B$	add matrices A and B
$A - B$	subtract matrix B from A
$A * B$	matrix multiplication of A and B
$A @ B$	element-wise multiplication of matrices A and B
A / B	element-wise division of matrix A by matrix B
$A == B$	element-wise equality evaluation between matrices A and B [caveat: use <code>approx_equal()</code> to test whether all corresponding elements are approximately equal]
$A != B$	element-wise non-equality evaluation between matrices A and B
$A >= B$	element-wise evaluation whether elements in matrix A are greater-than-or-equal to elements in B
$A <= B$	element-wise evaluation whether elements in matrix A are less-than-or-equal to elements in B
$A > B$	element-wise evaluation whether elements in matrix A are greater than elements in B
$A < B$	element-wise evaluation whether elements in matrix A are less than elements in B

Table 3: Subset of functions for matrix decompositions, factorisations, inverses and equation solvers.

Function	Description
<code>chol(X)</code>	Cholesky decomposition of symmetric positive-definite matrix X
<code>eig_sym(X)</code>	eigen decomposition of a symmetric/hermitian matrix X
<code>eig_gen(X)</code>	eigen decomposition of a general (non-symmetric/non-hermitian) square matrix X
<code>eig_pair(A, B)</code>	eigen decomposition for pair of general square matrices A and B
<code>inv(X)</code>	inverse of a square matrix X
<code>inv_sympd(X)</code>	inverse of symmetric positive definite matrix X
<code>lu(L, U, P, X)</code>	lower-upper decomposition of X , such that $PX = LU$ and $X = P'LU$
<code>null(X)</code>	orthonormal basis of the null space of matrix X
<code>orth(X)</code>	orthonormal basis of the range space of matrix X
<code>pinv(X)</code>	Moore-Penrose pseudo-inverse of a non-square matrix X
<code>qr(Q, R, X)</code>	QR decomposition of X , such that $QR = X$
<code>qr_econ(Q, R, X)</code>	economical QR decomposition
<code>qz(AA, BB, Q, Z, A, B)</code>	generalised Schur decomposition for pair of general square matrices A and B
<code>schur(X)</code>	Schur decomposition of square matrix X
<code>solve(A, B)</code>	solve a system of linear equations $AX = B$, where X is unknown
<code>svd(X)</code>	singular value decomposition of X
<code>svd_econ(X)</code>	economical singular value decomposition of X
<code>syl(X)</code>	Sylvester equation solver

Table 4: Subset of functions for generating vectors and matrices, showing their main form.

Function	Description
<code>linspace(start, end, n)</code>	vector with n elements, linearly spaced from $start$ upto (and including) end
<code>logspace(A, B, n)</code>	vector with n elements, logarithmically spaced from 10^A upto (and including) 10^B
<code>regspace(start, Δ, end)</code>	vector with regularly spaced elements: [$start, (start + \Delta), (start + 2\Delta), \dots, (start + M\Delta)$], where $M = \text{floor}((end - start)/\Delta)$, so that $(start + M\Delta) \leq end$
<code>randperm(n, m)</code>	vector with m elements, with random integers sampled without replacement from 0 to $n-1$
<code>randu(n_rows, n_cols)</code>	matrix with random values from a uniform distribution in the $[0,1]$ interval
<code>randn(n_rows, n_cols)</code>	matrix with random values from a normal/Gaussian distribution with $\mu = 0$ and $\sigma = 1$
<code>zeros(n_rows, n_cols)</code>	matrix with all elements set to zero
<code>ones(n_rows, n_cols)</code>	matrix with all elements set to one

Table 5: Subset of general functions of matrices, showing their main form. For functions with the *dim* argument, *dim = 0* indicates *traverse across rows* (ie. operate on all elements in a column), while *dim = 1* indicates *traverse across columns* (ie. operate on all elements in a row); by default *dim = 0*.

Function	Description
abs(A)	obtain element-wise magnitude of each element of matrix <i>A</i>
accu(A)	accumulate (sum) all elements of matrix <i>A</i> into a scalar
all(A,dim)	return a vector indicating whether all elements in each column or row of <i>A</i> are non-zero
any(A,dim)	return a vector indicating whether any element in each column or row of <i>A</i> is non-zero
approx_equal(A, B, met, tol)	return a bool indicating whether all corresponding elements in <i>A</i> and <i>B</i> are approx. equal
as_scalar(expression)	evaluate an expression that results in a 1×1 matrix, then convert the result to a pure scalar
clamp(A, min, max)	create a copy of matrix <i>A</i> with each element clamped to be between <i>min</i> and <i>max</i>
cond(A)	condition number of matrix <i>A</i> (the ratio of the largest singular value to the smallest)
conj(C)	complex conjugate of complex matrix <i>C</i>
cross(A, B)	cross product of <i>A</i> and <i>B</i> , assuming they are 3 dimensional vectors
cumprod(A, dim)	cumulative product of elements in each column or row of matrix <i>A</i>
cumsum(A, dim)	cumulative sum of elements in each column or row of matrix <i>A</i>
det(A)	determinant of square matrix <i>A</i>
diagmat(A, k)	interpret matrix <i>A</i> as a diagonal matrix (elements not on <i>k</i> -th diagonal are treated as zero)
diagvec(A, k)	extract the <i>k</i> -th diagonal from matrix <i>A</i> (default: <i>k = 0</i>)
diff(A, k, dim)	differences between elements in each column or each row of <i>A</i> ; <i>k</i> = number of recursions
dot(A,B)	dot product of <i>A</i> and <i>B</i> , assuming they are vectors with equal number of elements
eps(A)	distance of each element of <i>A</i> to next largest representable floating point number
expmat(A)	matrix exponential of square matrix <i>A</i>
find(A)	find indices of non-zero elements of <i>A</i> ; <i>find(A > k)</i> finds indices of elements greater than <i>k</i>
fliplr(A)	copy <i>A</i> with the order of the columns reversed
flipud(A)	copy <i>A</i> with the order of the rows reversed
imag(C)	extract the imaginary part of complex matrix <i>C</i>
ind2sub(size(A), index)	convert a linear index (or vector of indices) to subscript notation, using the size of matrix <i>A</i>
inplace_trans(A)	in-place / in-situ transpose of matrix <i>A</i>
join_rows(A, B)	append each row of <i>B</i> to its respective row of <i>A</i>
join_cols(A, B)	append each column of <i>B</i> to its respective column of <i>A</i>
kron(A, B)	Kronecker tensor product of <i>A</i> and <i>B</i>
log_det(x, sign, A)	log determinant of square matrix <i>A</i> , such that the determinant is $\exp(x) * \text{sign}$
logmat(A)	complex matrix logarithm of square matrix <i>A</i>
min(A, dim)	find the minimum in each column or row of matrix <i>A</i>
max(A, dim)	find the maximum in each column or row of matrix <i>A</i>
nonzeros(A)	return a column vector containing the non-zero values of matrix <i>A</i>
norm(A,p)	<i>p</i> -norm of matrix <i>A</i> , with <i>p</i> = 1, 2, ..., or <i>p</i> = "-inf", "inf", "fro"
normalise(A, p, dim)	return the normalised version of <i>A</i> , with each column or row normalised to unit <i>p</i> -norm
prod(A, dim)	product of elements in each column or row of matrix <i>A</i>
rank(A)	rank of matrix <i>A</i>
rcond(A)	estimate the reciprocal of the condition number of square matrix <i>A</i>
real(C)	extract the real part of complex matrix <i>C</i>
repmat(A, p, q)	replicate matrix <i>A</i> in a block-like fashion, resulting in <i>p</i> by <i>q</i> blocks of matrix <i>A</i>
reshape(A, r, c)	create matrix with <i>r</i> rows and <i>c</i> columns by copying elements from <i>A</i> column-wise
resize(A, r, c)	create matrix with <i>r</i> rows and <i>c</i> columns by copying elements and their layout from <i>A</i>
shift(A, n, dim)	copy matrix <i>A</i> with the elements shifted by <i>n</i> positions in each column or row
shuffle(A, dim)	copy matrix <i>A</i> with elements shuffled in each column or row
size(A)	obtain the dimensions of matrix <i>A</i>
sort(A, direction, dim)	copy <i>A</i> with elements sorted (in ascending or descending direction) in each column or row
sort_index(A, direction)	generate a vector of indices describing the sorted order of the elements in matrix <i>A</i>
sqrtmat(A)	complex square root of square matrix <i>A</i>
sum(A, dim)	sum of elements in each column or row of matrix <i>A</i>
sub2ind(size(A), row, col)	convert subscript notation (<i>row,col</i>) to a linear index, using the size of matrix <i>A</i>
symmatu(A) / symmatl(A)	generate symmetric matrix from square matrix <i>A</i>
strans(C)	simple matrix transpose of complex matrix <i>C</i> , without taking the conjugate
trans(A)	transpose of matrix <i>A</i> (for complex matrices, conjugate is taken); use <i>A.t()</i> for shorter form
trace(A)	sum of the elements on the main diagonal of matrix <i>A</i>
trapz(A, B, dim)	trapezoidal integral of <i>B</i> with respect to spacing in <i>A</i> , in each column or row of <i>B</i>
trimatu(A) / trimatl(A)	generate triangular matrix from square matrix <i>A</i>
unique(A)	return the unique elements of <i>A</i> , sorted in ascending order
vectorise(A, dim)	generate a column or row vector from matrix <i>A</i>

Table 6: Element-wise functions: matrix B is produced by applying a function to each element of matrix A .

Function	Description
exp(A)	base-e exponential: e^x
exp2(A)	base-2 exponential: 2^x
exp10(A)	base-10 exponential: 10^x
trunc_exp(A)	base-e exponential, truncated to avoid ∞
log(A)	natural log: $\log_e(x)$
log2(A)	base-2 log: $\log_2(x)$
log10(A)	base-10 log: $\log_{10}(x)$
trunc_log(A)	natural log, truncated to avoid $\pm\infty$
pow(A, p)	raise to the power of p: x^p
square(A)	square: x^2
sqrt(A)	square root: \sqrt{x}
floor(A)	largest integral value that is not greater than the input value
ceil(A)	smallest integral value that is not less than the input value
round(A)	round to nearest integer, with halfway cases rounded away from zero
trunc(A)	round to nearest integer, towards zero
erf(A)	error function
erfc(A)	complementary error function
lgamma(A)	natural log of the gamma function
sign(A)	signum function; for each element a in A , the corresponding element b in B is: $b = \begin{cases} -1 & \text{if } a < 0 \\ 0 & \text{if } a = 0 \\ +1 & \text{if } a > 0 \end{cases}$
trig(A)	trigonometric function, where <i>trig</i> is one of: <i>cos, acos, cosh, acosh, sin, asin, sinh, asinh, tan, atan, tanh, atanh</i>

Table 7: Subset of functions for statistics, showing their main form. For functions with the *dim* argument, *dim* = 0 indicates *traverse across rows* (ie. operate on all elements in a column), while *dim* = 1 indicates *traverse across columns* (ie. operate on all elements in a row); by default *dim* = 0.

Function/Class	Description
cor(A, B)	generate matrix of correlation coefficients between variables in A and B
cov(A, B)	generate matrix of covariances between variables in A and B
hist(A, centers, dim)	generate matrix of histogram counts for each column or row of A , using given bin centers
histc(A, edges, dim)	generate matrix of histogram counts for each column or row of A , using given bin edges
kmeans(means, A, k, ...)	cluster column vectors in matrix A into k disjoint sets, storing the set centers in <i>means</i>
principal(A)	principal component analysis of matrix A
running_stat	class for running statistics of a continuously sampled one dimensional signal
running_stat_vec	class for running statistics of a continuously sampled multi-dimensional signal
mean(A, dim)	find the mean in each column or row of matrix A
median(A, dim)	find the median in each column or row of matrix A
stddev(A, norm.type, dim)	find the standard deviation in each column or row of A , using specified normalisation
var(A, norm.type, dim)	find the variance in each column or row of matrix A , using specified normalisation

Table 8: Subset of functions for signal and image processing, showing their main form.

Function/Class	Description
conv(A, B)	1D convolution of vectors A and B
conv2(A, B)	2D convolution of matrices A and B
fft(A, n)	fast Fourier transform of vector A , with transform length n
ifft(C, n)	inverse fast Fourier transform of complex vector C , with transform length n
fft2(A, rows, cols)	fast Fourier transform of matrix A , with transform size of <i>rows</i> and <i>cols</i>
ifft2(C, rows, cols)	inverse fast Fourier transform of complex matrix C , with transform size of <i>rows</i> and <i>cols</i>
interp1(X, Y, XI, YI)	given a 1D function specified in vectors X (locations) and Y (values), generate vector YI containing interpolated values at given locations XI
interp2(X, Y, Z, XI, YI, ZI)	given a 2D function specified by matrix Z with coordinates given by vectors X and Y , generate matrix ZI containing interpolated values at coordinates given by vectors XI and YI

Table 9: Examples of Matlab/Octave syntax and conceptually corresponding PyArmadillo syntax. Note that for submatrix access the exact conversion from Matlab/Octave to PyArmadillo syntax will require taking into account that indexing starts at 0.

Matlab & Octave	PyArmadillo	Notes
A(1, 1) A(k, k)	A[0, 0] A[k-1, k-1]	indexing in PyArmadillo starts at 0, following Python convention
size(A,1) size(A,2) size(Q,3) numel(A)	A.n_rows A.n_cols Q.n_slices A.n_elem	member variables are read only Q is a cube (3D array) .n_elem indicates the total number of elements
A(:, k) A(k, :) A(:, p:q) A(p:q, :) A(p:q, r:s)	A[:, k] A[k, :] A[:, p:q] A[p:q, :] A[p:q, r:s]	read/write access to a specific column read/write access to a specific row read/write access to a submatrix spanning the specified cols read/write access to a submatrix spanning the specified rows
Q(p:q, r:s, t:u) Q(:, :, k) Q(:, :, k)	Q[p:q, r:s, t:u] Q[:, :, k] Q[single_slice, k]	Q is a cube (3D array) access one slice access one slice, with the slice provided as a matrix
A' A.'	A.t() or trans(A) A.st() or strans(A)	transpose (for complex matrices the conjugate is taken) simple transpose (for complex matrices the conjugate is not taken)
A = zeros(size(A)) A = ones(size(A)) A = zeros(k) A = ones(k)	A.zeros() A.ones() A = zeros(k, k) A = ones(k, k)	set all elements to zero set all elements to one create a matrix with elements set to zero create a matrix with elements set to one
C = complex(A,B)	C = cx_mat(A,B)	construct a complex matrix from two real matrices
A * B A .* B A ./ B A \ B A = A + 1 A = A - 1	A * B A @ B A / B solve(A,B) A += 1 A -= 1	* indicates matrix multiplication @ indicates element-wise multiplication / indicates element-wise division solve a system of linear equations
A = [1 2; 3 4;]	A = [[1, 2], [3, 4]]	
X = [A B] X = [A; B]	X = join_rows(A,B) X = join_cols(A,B)	
A save -ascii 'A.txt' A load -ascii 'A.txt'	A.print("A:") A.save("A.txt", raw_ascii) A.load("A.txt", raw_ascii)	print the entire contents of A Matlab/Octave matrices saved as ascii text are readable by PyArmadillo (and vice-versa)
A = rand(2,3) B = randn(4,5)	A = randu(2,3) B = randn(4,5)	randu generates uniformly distributed random numbers

Acknowledgements

We would like to thank our colleagues at Data61/CSIRO (Andrew Bolt, Dan Pagendam, Piotr Szul) for providing feedback and initial testing.

References

- [1] D. Abrahams and R. W. Grosse-Kunstleve. Building hybrid systems with Boost.Python. *C/C++ Users Journal*, 21(7), 2003.
- [2] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*. SIAM, 1999.
- [3] Apache Software Foundation. Apache license 2.0, 2004. <http://www.apache.org/licenses/LICENSE-2.0>.
- [4] D. Eddelbuettel and C. Sanderson. RcppArmadillo: Accelerating R with high-performance C++ linear algebra. *Computational Statistics & Data Analysis*, 71:1054–1063, 2014.
- [5] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, et al. Array programming with NumPy. *Nature*, 585:357–362, 2020.
- [6] Intel. Math Kernel Library (MKL), 2016. <http://software.intel.com/en-us/intel-mkl/>.
- [7] W. Jakob, J. Rhinelander, and D. Moldovan. pybind11 – seamless operability between C++11 and Python, 2017. <https://github.com/pybind/pybind11>.
- [8] S. Linge and H. P. Langtangen. *Programming for Computations - MATLAB/Octave*. Springer, 2016.
- [9] C. Sanderson and R. Curtin. Armadillo: a template-based C++ library for linear algebra. *Journal of Open Source Software*, 1:26, 2016.
- [10] C. Sanderson and R. Curtin. A user-friendly hybrid sparse matrix class in C++. In *Lecture Notes in Computer Science (LNCS)*, Vol. 10931, pages 422–430, 2018.
- [11] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17:261–272, 2020.
- [12] Z. Xianyi, W. Qian, and W. Saar. OpenBLAS: An optimized BLAS library, 2016. <http://www.openblas.net/>.