Armadillo: An Efficient Framework for Numerical Linear Algebra

Conrad Sanderson [†‡] and Ryan Curtin [◇]

[†] Data61 / CSIRO, Australia; [‡] Griffith University, Australia; [◇] NumFOCUS Inc., USA

Abstract—**A major challenge in the deployment of scientific software solutions is the adaptation of research prototypes to production-grade code. While high-level languages like MATLAB are useful for rapid prototyping, they lack the resource efficiency required for scalable production applications, necessitating translation into lower level languages like C++. Further, for machine learning and signal processing applications, the underlying linear algebra primitives, generally provided by the standard BLAS and LAPACK libraries, are unwieldy and difficult to use, requiring manual memory management and other tedium. To address this challenge, the Armadillo C++ linear algebra library provides an intuitive interface for writing linear algebra expressions that are easily compiled into efficient production-grade implementations. We describe the expression optimisations we have implemented in Armadillo, exploiting template metaprogramming. We demonstrate that these optimisations result in considerable efficiency gains on a variety of benchmark linear algebra expressions.**

Index Terms—**numerical linear algebra, BLAS, LAPACK, automated mapping, metaprogramming, expression optimisation.**

I. Introduction

Deployment and productisation of various machine learning and signal processing algorithms often requires conversion of research code written in a high-level language (eg., Matlab [12]) into a lower level language such as C or C++, which is considerably more resource efficient [22]. Resource efficiency is an important concern: in datacenter environments, the efficiency of production code is directly connected to cost (power costs and/or cloud resource costs). In environments with constrained computational resources, such as robots, unmanned aerial vehicles and spacecraft, efficiency is especially important as prototype code may be entirely unable to run on the target device due to limited memory or computational power.

Many algorithms inherently rely on numerical linear algebra operations, which are typically provided by the well-tested industry standard BLAS and LAPACK toolkits [2], [4], and their high-performance drop-in substitutes like OpenBLAS [25], MKL [13] and AOCL [1]. However, converting arbitrary linear algebra expressions into an efficient sequence of well-matched calls to BLAS and LAPACK routines is non-trivial [3], [17]; manual conversion can be laborious and error-prone, and requires good understanding of the intricacies of BLAS and LAPACK, including various trade-offs across available routines and storage formats. A further downside of directly using BLAS/LAPACK routines is that the resultant source code is quite verbose, has little similarity to the original mathematical expressions, involves keeping track of many supporting variables, and requires manual memory management. Such aspects significantly reduce the readability of the source code, raise the risk of bugs, and increase the maintenance burden [15], [19].

To address the above issues in a coherent framework, we have implemented the Armadillo linear algebra library for C++ [18], which automatically optimises mathematical expressions (both at compile-time and run-time) and efficiently maps them to BLAS/LAPACK routines, all while providing a user-friendly Matlab-like programming interface directly in C++. Armadillo essentially acts as a high-level *domain specific language* [16] built on top of the host C++ language, allowing for resource efficient numerical linear algebra without the many pain points of low-level code. This enables rapid and low risk conversion of research code into production environments, and even permits direct prototyping of algorithms within C++.

As an expository demonstration of the reduced maintenance burden when using Armadillo, consider the matrix expression $c = A^{-1}b$ for matrix A and vectors b and c, which represents the solution to a system of linear equations. Using Armadillo, it can be implemented directly in C++ as a single readable and maintainable line of code:

```
vec c = inv(A) * b;
```

Naively mapped, the above code will result in subsequent calls to three LAPACK and BLAS functions[1]: xGETRF, xGETRI, and xGEMV. Each of those three functions has between 6 and 11 parameters and may require manual allocation of workspace memory. In addition to hiding the verbosity and associated burdens with calls to BLAS and LAPACK functions, Armadillo is also able to reinterpret the expression and perform a better mapping to more efficient BLAS/LAPACK functions, avoiding the explicit matrix inverse operation.

Armadillo employs two strategies for automatically optimising mathematical expressions, both aiming to reduce computational effort: **(i)** compile-time fusion of operations to reduce the need for temporary objects, **(ii)** mixture of compile-time detection of expressions and run-time analysis of matrix properties, with the aim of re-ordering and translating operations. Both strategies extensively use C++ template metaprogramming concepts [9], [23], where the compiler is induced to reason at compile-time to generate code tailored for each expression.

We continue the paper as follows. Section II overviews the techniques for compile-time and run-time expression optimisation. Section III provides an empirical evaluation demonstrating the speedups obtained from the optimisations. The salient points and avenues for further exploitation are summarised in Section IV.

---

[1]The first letter of each BLAS/LAPACK function is replaced with x to express a set of functions that differ only in the associated element type (*float*, *double*, ...). For example, xGEMM represents the SGEMM, DGEMM, CGEMM and ZGEMM functions.

II. Expression Optimisation via Metaprogramming

Template metaprogramming induces the C++ compiler to run special programs written in a subset of the C++ language. Such metaprograms are executed entirely at compile-time, and can be used to produce compiled code that is specialised for arbitrary object and element types [23].

Rather than directly and immediately evaluating each component of a mathematical expression, Armadillo exploits template metaprogramming via lightweight marker objects that hold references to matrices and data associated with specific operations. The marker objects are generated via user-accessible functions (such as addition and multiplication) and store the identifier of each operation as a custom *type* only visible to the C++ compiler, rather than an explicit value. The marker objects can be chained together, leading to the full description of an arbitrary mathematical expression to be visible to the C++ compiler as an elaborate type, comprised as a tree of operation types. The evaluation of the entire expression is automatically performed when it is assigned to a target matrix. This approach is known as *delayed evaluation* (also known as *lazy evaluation*), and is in contrast to the traditional *eager evaluation* and *greedy evaluation* approaches [24].

As an illustrative example, let us consider the expression

```
Z = 0.4 * X + 0.6 * Y
```

where X, Y and Z are pre-defined Mat objects, each holding a 100×100 matrix. In a traditional eager evaluation approach, the 0.4 * X operation would be evaluated first, storing the intermediate result in a temporary matrix T1. The 0.6 * Y operation would then result in a secondary temporary matrix T2. The temporary matrices T1 and T2 would then be added, finally storing the result in matrix Z. This approach for the evaluation of the entire expression is suboptimal and inefficient, as it requires time-consuming memory allocation for the two temporary matrices and three separate loops over the associated matrix elements.

The delayed evaluation approach implemented in Armadillo aims to address such inefficiencies. Through overloading the * operator function, the operation 0.4 * X is not evaluated directly, but is instead automatically converted to a lightweight templated marker object named Op<Mat, op_mul>, which holds a reference to the X object and a copy of the 0.4 scalar multiplier. The nomenclature Op<...> indicates that Op is a C++ template class, with the items (types) between '<' and '>' specifying template parameters. A similar Op marker object is automatically constructed for the 0.6 * Y operation. The + operator function is overloaded to accept Mat objects and arbitrary marker objects, generating a templated Glue marker object that holds references to the given objects. In this example, it chains the two generated Op objects, resulting in the Glue object having the following type:

```
Glue< Op<Mat, op_mul>, Op<Mat, op_mul>, glue_plus >
```

The expression evaluation mechanism in Armadillo is then automatically invoked through the = operator defined in the Mat object. The mechanism interprets (at compile-time) the nested types in the template parameters of the given Glue object and automatically generates compiled instructions equivalent to:

```
for(int i=0; i<N; ++i) { Z[i] = 0.4 * X[i] + 0.6 * Y[i]; }
```

where N is the number of elements in matrices X, Y and Z, with X[i] indicating the i-th element in matrix X. Apart from the lightweight Op and Glue marker objects (which are automatically generated and pre-allocated at compile-time), no other temporary objects are generated. Furthermore, only one loop over the elements is required, instead of three separate loops in the traditional eager evaluation approach.

As a further efficiency enhancement, modern C++ compilers exploit aggressive optimisation strategies that are able to remove lightweight scaffolding objects. This results in the compiler producing machine code where the temporary Op and Glue objects are optimised away, leaving only code absolutely necessary for the specialised loop, tailored for the given expression. Moreover, this loop can be automatically *vectorised* by the C++ compiler, where low-level Single-Instruction-Multiple-Data (SIMD) instructions are exploited to achieve higher throughput [21].

The expression evaluation mechanisms in Armadillo include safety checks, to ensure that only compatible sizes can be used for each given operation. For example, checking that two matrices to be added or multiplied have conforming dimensions.

For mathematical expressions involving element-wise operations that can be chained, the evaluation mechanism is able to handle an arbitrary number of components (eg., matrices) within the given expressions. Other expressions are handled through detecting specific template patterns, possibly embedded within longer expressions. For example, the expression inv(A) * b is translated to the following Glue template type:

```
Glue< Op<Mat, op_inv>, Vec, glue_times >
```

The above pattern is detected at compile-time, and is automatically translated as a call to the xGESV function in LAPACK, which solves a system of linear equations without the matrix inverse.

In general, expressions with matrix multiplication are typically translated as calls to the xGEMM and xGEMV functions in BLAS, which are in turn multi-threaded and hand optimised for specific CPU architectures in high-performance implementations such as OpenBLAS [25].

Expression patterns are not necessarily blindly mapped to loops or BLAS/LAPACK functions. Specific patterns are further analysed at run-time, by analysing the properties of the constituent matrices. For example, run-time analysis is used for detecting that in the expression $A \cdot A^T$, the matrix multiplication involves the same matrix and results in a symmetric matrix. Rather than mapping the expression to the xGEMM function by default, the more efficient xSYRK function can be used, which exploits the symmetry property.

Analysis of matrix properties is also exploited in the evaluation of matrix multiplication chains. For example, in the expression $A \cdot B \cdot C \cdot D$, each of the possible matrix pairs is examined. The pair which results in the smallest matrix is multiplied first, thereby reducing computational effort in subsequent matrix multiplications. As such, it is possible for the entire expression to be evaluated right-to-left (while respecting general non-commutativity of matrix multiplication), rather than the traditional left-to-right order.

III. Empirical Evaluation

We evaluate the following representative set of expressions to demonstrate some of the optimisations automatically attainable by the expression processing frameworks implemented[2] in the Armadillo library.

(1). $C = 0.4 \cdot A + 0.6 \cdot B$; this is an instance of a compound expression involving element-wise addition of matrices and element-wise multiplication of matrices by scalars. A naive implementation evaluates each component separately, generating temporary matrices for $0.4 \cdot A$ and $0.6 \cdot B$, followed by adding the temporary matrices. An optimised implementation is able to bypass the generation of the temporaries, combining scalar multiplication and element addition into one loop that can exploit high-performance SIMD instructions present in modern CPUs [21]. SIMD instructions such as AVX-512 allow efficient processing of chunks of data in one hit instead of individual elements [5].

(2). $C = A_{(:,1)} + B_{(2,:)}^T$; this expression involves element-wise addition of submatrices (accessing individual columns and rows) in conjunction with matrix transpose. The notation $A_{(:,1)}$ denotes the first column of A, while $B_{(2,:)}$ denotes the second row of B. A naive implementation explicitly extracts the column and row into temporary vectors, followed by applying a transpose operation that generates a further temporary vector, which is then used for element-wise addition. An optimised implementation bypasses the generation of all temporary vectors as well as the explicit transpose operation, and instead accesses the matrix elements directly, performing an implicit transpose where required.

(3). $C = \mathrm{diagmat}(A) \cdot B$; this expression demonstrates matrix multiplication where one of the matrices is converted to a diagonal matrix. The $\mathrm{diagmat}(A)$ function indicates that all elements not on the main diagonal of A are assumed to be zero. In a naive implementation, the $\mathrm{diagmat}(A)$ function extracts the diagonal from A, and places it a temporary matrix. The temporary matrix (which is assumed by default to be dense) is then multiplied with B through a call to the standard xGEMM function in BLAS. An optimised implementation omits generating the temporary, and instead performs a specialised matrix multiplication which exploits sparsity by assuming that only the diagonal elements of A are non-zero.

(4). $C = \mathrm{diagmat}(A \cdot B)$; in this expression the result of matrix multiplication is converted into a diagonal matrix. A naive implementation would blindly evaluate $A \cdot B$ via the xGEMM function in BLAS and store the result in a temporary matrix, followed by extracting the diagonal from the temporary and placing it in the final result matrix. An optimised implementation is able to determine that only the diagonal elements of the matrix multiplication are required, thereby omitting unnecessary computations and temporaries.

(5). $k = \mathrm{trace}(A \cdot B)$; this expression is similar to the preceding $\mathrm{diagmat}(A \cdot B)$ expression, with the main difference that the diagonal elements of $A \cdot B$ are summed into the scalar k. In a naive implementation full matrix multiplication is performed, while an optimised implementation performs a partial matrix multiplication to obtain only the diagonal elements.

(6). $E = A_{m \times m} \cdot B_{m \times \frac{m}{2}} \cdot C_{\frac{m}{2} \times \frac{m}{3}} \cdot D_{\frac{m}{3} \times \frac{m}{4}}$. this is an instance of chained matrix multiplication resulting in a matrix. Here the matrices are progressively decreasing in size. A naive implementation would evaluate each of the matrix products in the standard left-to-right manner, disregarding the wider context of the expression. An optimised implementation can examine the sizes of all possible matrix products within the expression, and determine that evaluating the products in a reversed order will save computational effort.

(7). $k = a^T \cdot \mathrm{diagmat}(B) \cdot c$; this is an example of chained matrix multiplication that results in a scalar value, where a and c are column vectors. A naive implementation computes each component separately (matrix transpose and generation of diagonal matrix) resulting in temporary matrices, and then performs matrix multiplication involving the temporaries. An optimised implementation can examine the expression and determine that only a single and straightforward element-wise multiply-and-sum loop is required over the underlying components, avoiding unnecessary computations and generation of temporaries. This type of expression optimisation is invoked in Armadillo via the `as_scalar()` function.

(8). $B = A \cdot A^T$; this expression is seemingly straightforward, involving a matrix being multiplied with its transposed version, resulting in a symmetric matrix. A naive implementation disregards this fact and blindly calculates the matrix product by treating the two components as separate matrices after an explicit transpose operation. A semi-optimised implementation can avoid the explicit transpose by appropriate mapping to the xGEMM function in BLAS. However, a fully optimised implementation can detect that the two matrices to be multiplied are the same, and map the expression to the more efficient xDSYRK function in BLAS, which exploits the symmetry aspect and avoids unnecessary computations.

(9). $C = A^{-1} \cdot b$; this expression indicates that a solution to a system of linear equations is *implicitly* sought. A naive implementation ignores the intent of the expression and calculates the inverse of matrix A followed by a matrix multiplication. Calculating the inverse is not only computationally inefficient, but also potentially numerically unstable. An optimised implementation can detect the intent of the expression and map it to the more appropriate xGESV function in LAPACK, which finds the solution through a more numerically stable algorithm [2].

(10). $C = \mathrm{solve}(A, b)$ where A is a tri-diagonal band matrix; this expression indicates that a solution to a system of linear equations is *explicitly* sought, with A having a special sparse structure. A naive implementation would disregard the structure. An optimised implementation can analyse the matrix and choose a more tailored solver function in LAPACK, thereby exploiting the sparse structure to avoid superfluous computations.

For each of the above expressions, the following multiple matrix sizes are used, ranging from small to large: { 100×100, 250×250, 500×500, 1000×1000 }. The evaluation is done on a machine with an AMD Ryzen 7640U x86-64 CPU running at 3.5 GHz. All source code was compiled with the GCC 14.2 C++ compiler. We also used the open-source OpenBLAS 0.3.26 library which provides optimised implementations of BLAS and LAPACK routines [25].

The results shown in Fig. 1 demonstrate that the optimised handling of expressions in Armadillo leads to considerable reduction in computational effort. Across the considered expressions, the reduction in wall-clock time is often over 50%, and in several cases it is over 90%.

Fig. 2 shows a simple Armadillo-based C++ program to demonstrate its intuitive programming syntax. Fig. 3 lists a trace of corresponding internal function calls, hiding from the user the complexity of calling BLAS and LAPACK functions.

---

[2]Documentation for all the functionality available in Armadillo is available at https://arma.sourceforge.net/docs.html

(1) expression: $C = 0.4 \cdot A + 0.6 \cdot B$			
matrix size	**naive**	**optimised**	**reduction**
100×100	5.51×10^{-6}	2.26×10^{-6}	59.04%
250×250	4.04×10^{-5}	1.66×10^{-5}	58.89%
500×500	1.87×10^{-4}	6.95×10^{-5}	62.87%
1000×1000	2.45×10^{-3}	7.85×10^{-4}	67.90%

(2) expression: $C = A_{(:,1)} + B_{(2,:)}^T$			
matrix size	**naive**	**optimised**	**reduction**
100×100	9.52×10^{-8}	3.38×10^{-8}	64.50%
250×250	2.94×10^{-7}	1.05×10^{-7}	64.43%
500×500	7.01×10^{-7}	3.37×10^{-7}	51.94%
1000×1000	1.30×10^{-6}	7.38×10^{-7}	43.07%

(3) expression: $C = \mathrm{diagmat}(A) \cdot B$			
matrix size	**naive**	**optimised**	**reduction**
100×100	3.84×10^{-5}	2.80×10^{-6}	92.70%
250×250	6.49×10^{-4}	2.81×10^{-5}	95.67%
500×500	5.01×10^{-3}	2.01×10^{-4}	95.99%
1000×1000	4.14×10^{-2}	1.49×10^{-3}	96.40%

(4) expression: $C = \mathrm{diagmat}(A \cdot B)$			
matrix size	**naive**	**optimised**	**reduction**
100×100	3.88×10^{-5}	4.86×10^{-6}	87.47%
250×250	6.51×10^{-4}	3.99×10^{-5}	93.87%
500×500	5.02×10^{-3}	1.75×10^{-4}	96.51%
1000×1000	4.11×10^{-2}	1.93×10^{-3}	95.31%

(5) expression: $k = \mathrm{trace}(A \cdot B)$			
matrix size	**naive**	**optimised**	**reduction**
100×100	3.73×10^{-5}	5.98×10^{-11}	99.99%
250×250	6.42×10^{-4}	6.62×10^{-11}	99.99%
500×500	4.93×10^{-3}	6.71×10^{-11}	99.99%
1000×1000	4.03×10^{-2}	2.71×10^{-10}	99.99%

(6) expression: $E = A_{m \times m} \cdot B_{m \times \frac{m}{2}} \cdot C_{\frac{m}{2} \times \frac{m}{3}} \cdot D_{\frac{m}{3} \times \frac{m}{4}}$			
matrix size	**naive**	**optimised**	**reduction**
100×100	1.20×10^{-5}	6.17×10^{-6}	48.53%
250×250	2.05×10^{-4}	1.02×10^{-4}	50.20%
500×500	1.54×10^{-3}	7.94×10^{-4}	48.34%
1000×1000	1.21×10^{-2}	6.02×10^{-3}	50.17%

(7) expression: $k = a^T \cdot \mathrm{diagmat}(B) \cdot c$			
matrix size	**naive**	**optimised**	**reduction**
100×100	1.85×10^{-6}	7.21×10^{-10}	99.96%
250×250	1.14×10^{-5}	8.54×10^{-10}	99.99%
500×500	4.77×10^{-5}	7.21×10^{-10}	99.99%
1000×1000	1.99×10^{-4}	7.24×10^{-10}	99.99%

(8) expression: $B = A \cdot A^T$			
matrix size	**naive**	**optimised**	**reduction**
100×100	3.97×10^{-5}	3.35×10^{-5}	15.59%
250×250	6.65×10^{-4}	3.78×10^{-4}	43.19%
500×500	5.07×10^{-3}	2.67×10^{-3}	47.41%
1000×1000	4.32×10^{-2}	2.21×10^{-2}	48.89%

(9) expression: $C = A^{-1} \cdot b$			
matrix size	**naive**	**optimised**	**reduction**
100×100	1.47×10^{-4}	5.45×10^{-5}	62.92%
250×250	1.46×10^{-3}	4.69×10^{-4}	67.91%
500×500	8.23×10^{-3}	2.79×10^{-3}	66.16%
1000×1000	5.33×10^{-2}	1.90×10^{-2}	64.34%

(10) expression: $C = \mathrm{solve}(A, b)$ where A is a tri-diagonal			
matrix size	**naive**	**optimised**	**reduction**
100×100	8.11×10^{-5}	2.16×10^{-5}	73.40%
250×250	6.19×10^{-4}	7.40×10^{-5}	88.04%
500×500	3.31×10^{-3}	2.06×10^{-4}	93.77%
1000×1000	2.13×10^{-2}	1.30×10^{-3}	93.91%

Fig. 1: Comparison of time taken (in seconds) for various matrix expressions, using naive (non-optimised) and automatically optimised implementations within the Armadillo linear algebra library. Average wall-clock time across 1000 runs is reported. Evaluations were performed on an AMD Ryzen 7640U CPU, running at 3.5 GHz. Code was compiled with the GCC 14.2 C++ compiler with the following flags: `-O3 -march=native`. OpenBLAS 0.3.26 was used for optimised implementations of BLAS and LAPACK routines [25].

```
01: #include <armadillo>
02:
03: using namespace arma;
04:
05: int main()
06:   {
07:   // generate random 100x100 matrix
08:   mat A(100, 100, fill::randu);
09:
10:   // generate random 100x1 vector
11:   vec b(100, fill::randu);
12:
13:   // solve for x in random symmetric system AA'x = b
14:   vec x = solve( A * A.t(), b );
15:
16:   x.print("x:");
17:
18:   return 0;
19:   }
```

Fig. 2: A simple Armadillo-based C++ program, solving a random symmetric system of linear equations.

```
Op<T1, op_type>::Op(T1&) [T1 = Mat; op_type = op_htrans]
operator*(T1&, T2&) [T1 = Mat; T2 = Op<Mat,op_htrans>]
Glue<T1, T2, glue_type>::Glue(T1&, T2&) [T1 = Mat; T2 = Op<Mat,op_htrans>; glue_type = glue_times]
solve(Base<double, T1>&, Base<double, T2>&)
Glue<T1, T2, glue_type>::Glue(T1&, T2&) [... glue_type = glue_solve_gen_def]
Col::Col(Base<double,T1>&) [T1 = Glue<Glue<Mat,Op<Mat,op_htrans>,glue_times>,Mat,glue_solve_gen_def>]
Mat::operator=(Glue<T1, T2, glue_type>&) [... glue_type = glue_solve_gen_def]
glue_solve_gen_def::apply(Mat&, Glue<T1, T2, glue_solve_gen_def>&)
glue_solve_gen_full::apply(Mat&, Base<double, T1>&, uword)
Mat::Mat(Glue<T1, T2, glue_type>&) [T1 = Mat; T2 = Op<Mat,op_htrans>; glue_type = glue_times]
glue_times::apply(Mat&, Glue<T1, T2, glue_times>&) [T1 = Mat; T2 = Op<Mat,op_htrans>]
glue_times_redirect<2>::apply(Mat&, Glue<T1, T2, glue_times>&) [T1 = Mat; T2 = Op<Mat,op_htrans>]
glue_times::apply(Mat&, TA&, TB&, double) [trans_A = false; trans_B = true; TA = Mat; TB = Mat]
Mat::set_size(uword, uword) [uword = long long unsigned int] [in_n_rows: 100; in_n_cols: 100]
Mat::init(): acquiring memory
blas::syrk(...)
glue_solve_gen_full::apply(): detected square system
band_helper::is_band(uword&, uword&, Mat&, uword) [uword = long long unsigned int]
trimat_helper::is_triu(Mat&)
trimat_helper::is_tril(Mat&)
glue_solve_gen_full::apply(): rcond + sym
auxlib::solve_sym_rcond(Mat&, double&, Mat&, Base<double, T1>&) [T1 = Mat; ...]
Mat::operator=(Mat&) [this: e0a67920; in_mat: e0a67860]
Mat::init_warm(uword, uword) [uword = long long unsigned int] [in_n_rows: 100; in_n_cols: 1]
Mat::init(): acquiring memory
lapack::lansy(...)
lapack::sytrf(...)
lapack::sytrs(...)
lapack::sycon(...)
Mat::destructor: releasing memory
```

Fig. 3: An abridged trace of internal function calls and debugging messages resulting from line 14 in Fig. 2, containing the expression `vec x = solve( A * A.t(), b )`.

IV. Conclusion

As we have seen, Armadillo facilitates easy and maintainable representation of arbitrary linear algebra expressions in C++ that are efficiently mapped to underlying BLAS and LAPACK operations. Users do not need to worry about cumbersome manual memory management or complicated calls to BLAS and LAPACK subroutines. There is virtually no performance penalty for the abstractions provided by Armadillo. Moreover, through under-the-hood template metaprogramming and automatic optimisations of expressions, Armadillo can achieve considerable reductions in processing time over direct and/or naive implementations.

Work on Armadillo started in 2008. Over the years the library has been iteratively and collaboratively developed with feedback from the wider scientific and engineering communities. The library provides over 200 functions; in addition to elementary operations, there are functions for statistics, signal processing, non-contiguous submatrix views, and various matrix factorisations. The library is currently comprised of about 135,000 lines of templated code, excluding BLAS and LAPACK routines. Support is provided for matrices with single- and double-precision floating point elements (in both real and complex forms), as well as integer elements. Dense and sparse storage formats are supported.

Armadillo is now in a mature state and in wide production use. For example, Armadillo has been successfully used to accelerate computations in open-source projects such as the *ensmallen* library for numerical optimisation [6] and the *mlpack* library for machine learning [8], which provide production-ready applications for a variety of environments, including low-resource devices such as small microcontrollers. Armadillo has also been used for accelerating over 1000 packages for the R statistical environment [10].

Armadillo can be obtained from https://arma.sourceforge.net, with the source code provided under the permissive Apache 2.0 license [14], [20], which allows unencumbered use in commercial products. Armadillo is also included as part of all major Linux distributions.

In future work we plan to extend Armadillo to include the support for half-precision floating point and 'brain floating point' (BF16) element types [11], as well as to bring the same kinds of expression optimisations to GPU-based linear algebra via the companion Bandicoot library [7].

Acknowledgements

References

[1] Advanced Micro Devices. AMD Optimizing CPU Libraries (AOCL), 2024. https://www.amd.com/en/developer/aocl.html.

[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] D. Berényi, A. Leitereg, and G. Lehel. Towards scalable pattern-based optimization for dense linear algebra. *Concurrency and Computation: Practice and Experience*, 30(22), 2018.

[4] L. S. Blackford, J. Demmel, J. Dongarra, I. Duff, S. Hammarling, et al. An updated set of basic linear algebra subprograms (BLAS). *ACM Transactions on Mathematical Software*, 28(2):135–151, 2002.

[5] J. M. Cebrian, L. Natvig, and M. Jahre. Scalability analysis of AVX-512 extensions. *The Journal of Supercomputing*, 76(3):2082–2097, 2020.

[6] R. R. Curtin, M. Edel, R. G. Prabhu, S. Basak, Z. Lou, and C. Sanderson. The ensmallen library for flexible numerical optimization. *Journal of Machine Learning Research*, 22(166), 2021.

[7] R. R. Curtin, M. Edel, and C. Sanderson. Bandicoot: C++ library for GPU linear algebra and scientific computing. *arXiv:2308.03120*, 2023.

[8] R. R. Curtin, M. Edel, O. Shrit, et al. mlpack 4: a fast, header-only C++ machine learning library. *Journal of Open Source Software*, 8(82), 2023.

[9] K. Czarnecki, U. Eisenecker, R. Glück, D. Vandevoorde, and T. Veldhuizen. Generative programming and active libraries. In *Lecture Notes in Computer Science (LNCS)*, volume 1766, pages 25–39, 2000.

[10] D. Eddelbuettel and C. Sanderson. RcppArmadillo: Accelerating R with high-performance C++ linear algebra. *Computational Statistics and Data Analysis*, 71:1054–1063, 2014.

[11] G. Henry, P. T. P. Tang, and A. Heinecke. Leveraging the bfloat16 artificial intelligence datatype for higher-precision computations. In *IEEE Symposium on Computer Arithmetic (ARITH)*, pages 69–76, 2019.

[12] D. J. Higham and N. J. Higham. *MATLAB Guide*. SIAM, 3rd edition, 2017.

[13] Intel Corporation. Math Kernel Library (MKL), 2025. https://software.intel.com/mkl.

[14] X. Li, Y. Zhang, C. Osborne, et al. Systematic literature review of commercial participation in open source software. *ACM Transactions on Software Engineering and Methodology*, 34(2), 2025.

[15] R. Malhotra and A. Chug. Software maintainability: Systematic literature review and current trends. *International Journal of Software Engineering and Knowledge Engineering*, 26(8):1221–1253, 2016.

[16] M. Mernik, J. Heering, and A. M. Sloane. When and how to develop domain-specific languages. *ACM Computing Surveys*, 37(4):316–344, 2005.

[17] C. Psarras, H. Barthels, and P. Bientinesi. The linear algebra mapping problem: Current state of linear algebra languages and libraries. *ACM Transactions on Mathematical Software*, 48(3), 2022.

[18] C. Sanderson and R. Curtin. Armadillo: a template-based C++ library for linear algebra. *Journal of Open Source Software*, 1(2):26, 2016.

[19] H. M. Sneed. A cost model for software maintenance & evolution. In *IEEE International Conference on Software Maintenance*, 2004.

[20] A. St. Laurent. *Understanding Open Source and Free Software Licensing*. O'Reilly Media, 2004.

[21] K. Stock, L.-N. Pouchet, and P. Sadayappan. Using machine learning to improve automatic vectorization. *ACM Transactions on Architecture and Code Optimization*, 8(4), 2012.

[22] B. Stroustrup. *Programming: Principles and Practice Using C++*. Addison-Wesley, 3rd edition, 2024.

[23] D. Vandevoorde, N. Josuttis, and D. Gregor. *C++ Templates: The Complete Guide*. Addison-Wesley, 2nd edition, 2017.

[24] D. A. Watt. *Programming Language Design Concepts*. Wiley, 2004.

[25] Z. Xianyi and M. Kroeker. OpenBLAS: An optimized BLAS library, 2025. http://www.openmathlib.org/OpenBLAS/.