

Practical Sparse Matrices in C++ with Hybrid Storage and Template-Based Expression Optimisation

Conrad Sanderson

conradsand@ieee.org
Data61, CSIRO, Australia
University of Queensland, Australia
Arroyo Consortium

Ryan Curtin

ryan@ratml.org
RelationalAI, USA
Arroyo Consortium

Abstract

Despite the importance of sparse matrices in numerous fields of science, software implementations remain difficult to use for non-expert users, generally requiring the understanding of underlying details of the chosen sparse matrix storage format. In addition, to achieve good performance, several formats may need to be used in one program, requiring explicit selection and conversion between the formats. This can be both tedious and error-prone, especially for non-expert users. Motivated by these issues, we present a user-friendly and open-source sparse matrix class for the C++ language, with a high-level application programming interface deliberately similar to the widely used MATLAB language. This facilitates prototyping directly in C++ and aids the conversion of research code into production environments. The class internally uses two main approaches to achieve efficient execution: (i) a hybrid storage framework, which automatically and seamlessly switches between three underlying storage formats (compressed sparse column, Red-Black tree, coordinate list) depending on which format is best suited and/or available for specific operations, and (ii) a template-based meta-programming framework to automatically detect and optimise execution of common expression patterns. Empirical evaluations on large sparse matrices with various densities of non-zero elements demonstrate the advantages of the hybrid storage framework and the expression optimisation mechanism.

- **Keywords:** mathematical software; C++ language; sparse matrix; numerical linear algebra
- **AMS MSC2010 codes:** 68N99; 65Y04; 65Y15; 65F50
- **Associated source code:** <http://arma.sourceforge.net>
- Published as:

Conrad Sanderson and Ryan Curtin.

Practical Sparse Matrices in C++ with Hybrid Storage and Template-Based Expression Optimisation.

Mathematical and Computational Applications, Vol. 24, No. 3, 2019.

<https://doi.org/10.3390/mca24030070>

1 Introduction

Recent decades have seen the frontiers of scientific computing increasingly push towards the use of larger and larger datasets. In fact, frequently the data to be represented is so large that it cannot fully fit into working memory. Fortunately, in many cases the data has many zeros and can be represented in a compact manner, allowing users to work with sparse matrices of extreme size with few non-zero elements. However, converting code from using dense matrices to using sparse matrices, a common task when scaling code to larger data, is not always straightforward.

Current open-source frameworks may provide several separate sparse matrix classes, each with its own data storage format. For example, SciPy [25] has 7 sparse matrix classes, where each storage format is best suited for efficient execution of a specific set of operations (eg., incremental matrix construction vs. matrix multiplication). Other frameworks may provide only one sparse matrix class, with considerable runtime penalties if it is not used in the right way. This can be challenging and bewildering for users who simply want to create and use sparse matrices, and do not have the time, expertise, nor desire to understand the advantages and disadvantages of each format. To achieve good performance, several formats may need to be used in one program, requiring explicit selection and conversion between the formats. This multitude of sparse matrix classes complicates the programming task, adds to the maintenance burden, and increases the likelihood of bugs.

Driven by the above concerns, we have devised a practical and user-friendly sparse matrix class for the C++ language [31]. The sparse matrix class uses a hybrid storage framework, which *automatically* and *seamlessly* switches between three data storage formats, depending on which format is best suited and/or available for specific operations:

- Compressed Sparse Column (CSC), used for efficient and nuanced implementation of core arithmetic operations such as matrix multiplication and addition, as well as efficient reading of individual elements;
- Red-Black Tree (RBT), used for both robust and efficient incremental construction of sparse matrices (i.e., construction via setting individual elements one-by-one, not necessarily in order);
- Coordinate List (COO), used for low-maintenance and straightforward implementation of relatively complex and/or lesser-used sparse matrix functionality.

The COO format is important to point out, as the source code for the sparse matrix class is distributed and maintained as part of the open-source Armadillo library (arma.sourceforge.net) [27]. Due to its simpler nature, the COO format facilitates functionality contributions from time-constrained and/or non-expert users, as well as reducing maintenance and debugging overhead for the library maintainers.

To further promote efficient execution, the sparse matrix class internally implements a delayed evaluation framework [21] based on template meta-programming [1, 33] combined with operator overloading [31]. In delayed evaluation, the evaluation of a given compound mathematical expression is delayed until its value is required (i.e., assigned to a variable). This is in contrast to eager evaluation (also known as strict evaluation), where each component of a compound expression is evaluated immediately. As such, the delayed evaluation framework allows automatic compile-time analysis of compound expressions, which in turns allows for automatic detection and optimisation of common expression patterns. For example, several operations can be combined to reduce the required computational effort.

Overall, the sparse matrix class and its associated functions provide a high-level application programming interface (function syntax) that is intuitive, close to a typical dense matrix interface, and deliberately similar to MATLAB. This can help with rapid transition of dense-specific code to sparse-specific code, facilitates prototyping directly in C++, and aids the conversion of research code into production environments.

While there are many other sparse matrix implementations in existence, to our knowledge the presented approach is the first to offer a unified interface with automatic format switching under the hood. Most toolkits are limited to either a single format or multiple formats the user must manually convert between. The comprehensive SPARSKIT package [26] contains 16, and SciPy contains seven formats [25]. In these toolkits the user must manually convert between formats. On the other hand, both MATLAB and GNU Octave [12] contain sparse matrix implementations, but they supply only the CSC format [10], meaning that users must write their code in special ways to ensure its efficiency [22]. This is a similar situation to the Blaze library (bitbucket.org/blaze-lib/blaze) [17], which uses a compressed format with either column- or row-major orientation. Users are explicitly discouraged from individual element insertions and, for efficiency, must construct their sparse matrices in a restricted and cumbersome environment of batch insertion. Furthermore, simply reading elements from the matrix via standard row and column indexing can result in temporary insertion of elements into the matrix. The Eigen C++ matrix library (eigen.tuxfamily.org) uses a specialised sparse matrix format which has deliberate redundancy and overprovisioned storage. While this can help with reducing the computational effort of element insertion in some situations, it requires manual care to maintain storage efficiency. Furthermore, as the cost of random insertion of elements is still high, the associated documentation recommends to manually construct a COO-like representation of all the elements, from which the actual sparse matrix is then constructed. The IT++ library (itpp.sourceforge.net) has a cumbersome sparse matrix class with a custom compressed column format that also employs overprovisioned storage. The format is less efficient storage-wise than CSC unless explicit manual care is taken. Data is stored in unordered fashion which allows for faster element insertion than CSC, but at the cost of reduced performance for linear algebra operations. Overall, the landscape of sparse matrix implementations is composed of libraries where the user must be aware of some of the internal storage details and the associated maintenance in order to produce efficient code; this is not ideal.

To make the situation even more complex, there are also numerous other sparse matrix formats [26, 11, 3]. Examples include the modified compressed row/column format (intended for sparse matrices where the main diagonal has all non-zero elements), block compressed storage format (intended for sparse matrices

with dense submatrices), diagonal format (intended for straightforward storage of banded sparse matrices under the assumption of constant bandwidth), and the various skyline formats (intended for more efficient storage of banded sparse matrices with irregular bandwidth). As these formats are focused on specialised use cases, their utility is typically not very general. Thus we have currently opted against including these formats in our hybrid framework, though it would be relatively easy to accommodate more formats in the future.

The paper is continued as follows. In Section 2 we overview the functionality provided by the sparse matrix class and its associated functions. The delayed evaluation approach is overviewed in Section 3. In Section 4 we describe the underlying storage formats used by the class, and the scenarios that each of the formats is best suited for. In Section 5 we discuss the costs for switching between the formats. Section 6 provides an empirical evaluation showing the advantages of the hybrid storage framework and the delayed evaluation approach. The salient points and avenues for further exploitation are summarised in Section 7. This article is a thoroughly revised and extended version of our earlier work [28].

2 Functionality

The sparse matrix class and its associated functions provide a user-friendly suite of essential sparse linear algebra functionality, including fundamental operations such as addition, matrix multiplication, and submatrix manipulation. The class supports storing elements as integers, single- and double-precision floating point numbers, as well as complex numbers. Various sparse eigendecompositions and linear equation solvers are provided through integration with low-level routines in the de-facto standard ARPACK [19] and SuperLU libraries [20]. The resultant high-level functions automatically take care of tedious and cumbersome details such as memory management, allowing users to concentrate their programming effort on mathematical details.

C++ language features such as overloading of operators (eg., * and +) [31] are exploited to allow mathematical operations with matrices to be expressed in a concise and easy-to-read manner, in a similar fashion to the proprietary MATLAB language. For example, given sparse matrices A, B, and C, a mathematical expression such as

$$D = \frac{1}{2}(A + B) \cdot C^T$$

can be written directly in C++ as

```
sp_mat D = 0.5 * (A + B) * C.t();
```

where `sp_mat` is our sparse matrix class. Figure 1 contains a complete C++ program which briefly demonstrates usage of the sparse matrix class, while Table 1 lists a subset of the available functionality.

The aggregate of the sparse matrix class, operator overloading, and associated functions on sparse matrices is an instance of a Domain Specific Language (sparse linear algebra in this case) embedded within the host C++ language [23, 29]. This allows complex algorithms relying on sparse matrices to be easily developed and integrated within a larger C++ program, making the sparse matrix class directly useful in application/product development.

3 Template-Based Optimisation of Compound Expressions

The sparse matrix class uses a delayed evaluation approach, allowing several operations to be combined to reduce the amount of computation and/or temporary objects. In contrast to brute-force evaluations, delayed evaluation can provide considerable performance improvements as well as reduced memory usage [34]. The delayed evaluation machinery is implemented through template meta-programming [1, 33], where a type-based signature of a compound expression (set of consecutive mathematical operations) is automatically constructed. The C++ compiler is then automatically induced to detect common expression patterns at compile time, followed by selecting the most computationally efficient implementations.

As an example of the possible efficiency gains, let us consider the expression $\text{trace}(A.t() * B)$, which often appears as a fundamental quantity in semidefinite programs [32]. These computations are thus used in a wide variety of diverse fields, most notably machine learning [6, 14, 18]. A brute-force implementation would evaluate the transpose first, $A.t()$, and store the result in a temporary matrix T_1 . The next operation would be a time consuming matrix multiplication, $T_1 * B$, with the result stored in another temporary matrix T_2 . The trace operation (sum of diagonal elements) would then be applied on T_2 . The explicit transpose, full

```

#include <armadillo>
using namespace arma;

int main()
{
  // generate random sparse 1000x1000 matrix with 1% density of non-zero values,
  // with uniform distribution of values in the [0,1] interval
  sp_mat A = sprandu(1000, 1000, 0.01);

  // multiply A by its transpose
  sp_mat B = A * A.t();

  // add scalar to main diagonal
  B.diag() += 0.1;

  // declare dense vector and matrix
  vec eigvals; mat eigvecs;

  // find 3 eigenvectors of sparse matrix B
  eigs_sym(eigvals, eigvecs, B, 3);

  return 0;
}

```

Figure 1: A small C++ program to demonstrate usage of the sparse matrix class (`sp_mat`).

Function	Description
<code>sp_mat X(1000,2000)</code>	Declare sparse matrix with 1000 rows and 2000 columns
<code>sp_cx_mat X(1000,2000)</code>	As above, but use complex elements
<code>X(1,2) = 3</code>	Assign value 3 to element at location (1,2) of matrix X
<code>X = 4.56 * A</code>	Multiply matrix A by scalar
<code>X = A + B</code>	Add matrices A and B
<code>X = A * B</code>	Multiply matrices A and B
<code>X(span(1,2), span(3,4))</code>	Provide read/write access to submatrix of X
<code>X.diag(k)</code>	Provide read/write access to diagonal k of X
<code>X.print()</code>	Print matrix X to terminal
<code>X.save(filename, format)</code>	Store matrix X as a file
<code>speye(rows, cols)</code>	Generate sparse matrix with values on the main diagonal set to one
<code>sprandu(rows, cols, density)</code>	Generate sparse matrix with random non-zero elements
<code>sum(X, dim)</code>	Sum of elements in each column ($dim=0$) or row ($dim=1$)
<code>min(X, dim); max(X, dim)</code>	Obtain extremum value in each column ($dim=0$) or row ($dim=1$)
<code>X.t()</code> or <code>trans(X)</code>	Return transpose of matrix X
<code>kron(A, B)</code>	Kronecker tensor product of matrices A and B
<code>repmat(X, rows, cols)</code>	Replicate matrix X in block-like fashion
<code>norm(X, p)</code>	Compute p -norm of vector or matrix X
<code>normalise(X, p, dim)</code>	Normalise each column ($dim=0$) or row ($dim=1$) to unit p -norm
<code>trace(A.t() * B)</code>	Compute trace of $A^T B$ without explicit transpose and multiplication
<code>diagmat(A + B)</code>	Obtain diagonal matrix from $A + B$ without full matrix addition
<code>eigs_gen(eigval, eigvec, X, k)</code>	Compute k largest eigenvalues and eigenvectors of matrix X
<code>svds(U, s, V, X, k)</code>	Compute k singular values and singular vectors of matrix X
<code>x = spsolve(A, b)</code>	Solve sparse system $Ax = b$ for x

Table 1: Subset of available functionality for the sparse matrix class, with brief descriptions. Optional additional arguments have been omitted for brevity. See <http://arma.sourceforge.net/docs.html> for more detailed documentation.

matrix multiplication and creation of the temporary matrices is suboptimal from an efficiency point of view, as for the trace operation we require only the diagonal elements of the $A.t() * B$ expression.

Template-based expression optimisation can avoid the unnecessary operations. Let us declare two lightweight objects, `Op` and `Glue`, where `Op` objects are used for representing unary operations, while `Glue` objects are used for representing binary operations. The objects are lightweight as they do not store actual sparse matrix data; instead the objects only store references to matrices and/or other `Op` and `Glue` objects. Ternary and more complex operations are represented through combinations of `Op` and `Glue` objects. The exact type of each `Op` and `Glue` object is automatically inferred from a given mathematical expression through template meta-programming.

In our example, the expression $A.t()$ is automatically converted to an instance of the lightweight `Op` object with the following type:

```
Op<sp_mat, op_trans>
```

where `Op<...>` indicates that `Op` is a template class, with the items between '`<`' and '`>`' specifying template parameters. In this case the `Op<sp_mat, op_trans>` object type indicates that a reference to a matrix is stored and that a transpose operation is requested. In turn, the compound expression $A.t() * B$ is converted to an instance of the lightweight `Glue` object with the following type:

```
Glue< Op<sp_mat, op_trans>, sp_mat, glue_times>
```

where the `Glue` object type in this case indicates that a reference to the preceding `Op` object is stored, a reference to a matrix is stored, and that a matrix multiplication operation is requested. In other words, when a user writes the expression $\text{trace}(A.t() * B)$, the C++ compiler is induced to represent it internally as $\text{trace}(\text{Glue}\langle \text{Op}\langle \text{sp_mat}, \text{op_trans}\rangle, \text{sp_mat}, \text{glue_times}\rangle(A,B))$.

There are several implemented forms of the `trace()` function, one of which is automatically chosen by the C++ compiler to handle the `Glue< Op<sp_mat, op_trans>, sp_mat, glue_times>` expression. The specific form of `trace()` takes references to the A and B matrices, and executes a *partial* matrix multiplication to obtain only the diagonal elements of the $A.t() * B$ expression. All of this is accomplished without generating temporary matrices. Furthermore, as the `Glue` and `Op` objects only hold references, they are in effect optimised away by modern C++ compilers [33]: the resultant machine code appears as if the `Glue` and `Op` objects never existed in the first place.

The template-based delayed evaluation approach has also been employed for other functions, such as the `diagmat()` function, which obtains a diagonal matrix from a given expression. For example, in the expression $\text{diagmat}(A + B)$, only the diagonal components of the $A + B$ expression are evaluated.

4 Storage Formats for Sparse Data

We have chosen the three underlying storage formats (CSC, RBT, COO) to give overall efficient execution across several use cases, as well as to minimise the difficulty of implementation and code maintenance burden where possible. Specifically, our focus is on the following main use cases:

1. Flexible ad-hoc construction and element-wise modification of sparse matrices via unordered insertion of elements, where each new element is inserted at a random location.
2. Incremental construction of sparse matrices via quasi-ordered insertion of elements, where each new element is inserted at a location that is past all the previous elements according to column-major ordering.
3. Multiplication of dense vectors with sparse matrices.
4. Multiplication of two sparse matrices.
5. Operations involving bulk coordinate transformations, such as flipping matrices column- or row-wise.

The three storage formats as well as their benefits and limitations are briefly described below. We use N to indicate the number of non-zero elements of the matrix, while n_rows and n_cols indicate the number of rows and columns, respectively. Examples of the formats are shown in Figure 2.

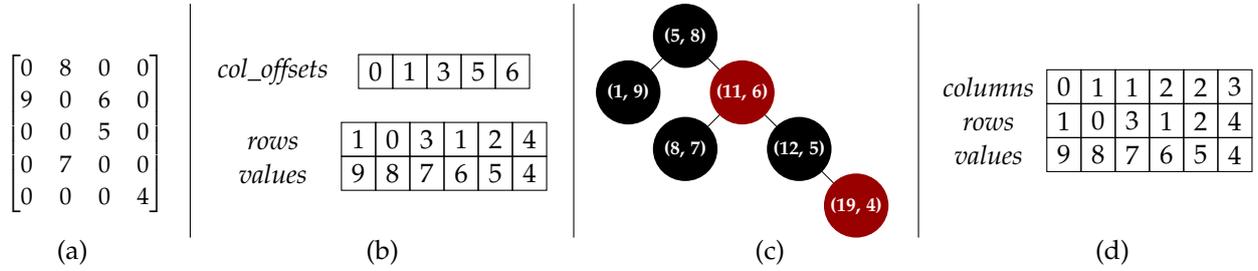


Figure 2: Illustration of sparse matrix representations: (a) example sparse matrix with 5 rows, 4 columns and 6 non-zero values, shown in traditional mathematical notation; (b) corresponding CSC representation; (c) corresponding RBT representation, where each node is expressed by (i, v) , with i indicating a linearly encoded matrix location and v indicating the value held at that location; (d) corresponding COO representation. Following C++ convention [31], we use zero-based indexing.

4.1 Compressed Sparse Column (CSC)

The CSC format [26] uses column-major ordering where the elements are stored column-by-column, with consecutive non-zero elements in each column stored consecutively in memory. Three arrays are used to represent a sparse matrix:

1. The *values* array, which is a contiguous array of N floating point numbers holding the non-zero elements.
2. The *rows* array, which is a contiguous array of N integers holding the corresponding row indices (ie., the n -th entry contains the row of the n -th element).
3. The *col_offsets* array, which is a contiguous array of $n_cols + 1$ integers holding offsets to the *values* array, with each offset indicating the start of elements belonging to each column.

Following C++ convention [31], all arrays use zero-based indexing, ie., the initial position in each array is denoted by 0. For consistency, element locations within a matrix are also encoded as starting at zero, ie., the initial row and column are both denoted by 0. Furthermore, the row indices for elements in each column are kept sorted in ascending manner. In many applications, sparse matrices have more non-zero elements than the number of columns, leading to the *col_offsets* array being typically much smaller than the *values* array.

Let us denote the i -th entry in the *col_offsets* array as $c[i]$, the j -th entry in the *rows* array as $r[j]$, and the n -th entry in the *values* array as $v[n]$. The number of non-zero elements in column i is determined using $c[i+1] - c[i]$, where, by definition, $c[0]$ is always 0 and $c[n_cols]$ is equal to N . If column i has non-zero elements, then the first element is obtained via $v[c[i]]$, and $r[c[i]]$ is the corresponding row of the element. An example of this format is shown in Figure 2(b).

The CSC format is well-suited for efficient sparse linear algebra operations such as vector-matrix multiplication. This is due to consecutive non-zero elements in each column being stored next to each other in memory, which allows modern CPUs to speculatively read ahead elements from the main memory into fast cache memory [24]. The CSC format is also suited for operations that do not change the structure of the matrix, such as element-wise operations on the non-zero elements (eg., multiplication by a scalar). The format also affords relatively efficient random element access; to locate an element (or determine that it is not stored), a single lookup to the beginning of the desired column can be performed, followed by a binary search [7] through the *rows* array to find the element.

While the CSC format provides a compact representation yielding efficient execution of linear algebra operations, it has two main disadvantages. The first disadvantage is that the design and implementation of efficient algorithms for many sparse matrix operations (such as matrix-matrix multiplication) can be non-trivial [4, 26]. This stems not only from the sparse nature of the data, but also due to the need to (i) explicitly keep track of the column offsets, (ii) ensure that the row indices for elements in each column are sorted in ascending manner, and (iii) ensure that any zeros resulting from an operation are not stored. In our experience, designing and implementing new and efficient matrix processing functions directly in the CSC format (which do not have prior publicly available implementations) can be time-consuming and prone to subtle bugs.

The second disadvantage of CSC is the computational effort required to insert a new element [10]. In the worst-case scenario, memory for three new larger-sized arrays (containing the values and locations) must first be allocated, the position of the new element determined within the arrays, data from the old arrays copied to the new arrays, data for the new element placed in the new arrays, and finally, the memory used by the old arrays deallocated. As the number of elements in the matrix grows, the entire process becomes slower.

There are opportunities for some optimisation, such as using oversized storage to reduce memory allocations, where a new element past all the previous elements can be readily inserted. However, this does not help when a new non-zero element is inserted between two existing non-zero elements. It is also possible to perform batch insertions with some speedup by first sorting all the elements to be inserted and then merging with the existing data arrays.

The CSC format was chosen over the related Compressed Sparse Row (CSR) format [26] for two main reasons: (i) to ensure compatibility with external libraries such as the SuperLU solver [20], and (ii) to ensure consistency with the surrounding infrastructure provided by the Armadillo library, which uses column-major dense matrix representation to take advantage of low-level functions provided by LAPACK [2].

4.2 Red-Black Tree (RBT)

To address the efficiency problems with element insertion at arbitrary locations, we first represent each element as a 2-tuple, $l = (index, value)$, where *index* encodes the location of the element as $index = row + column \times n_rows$. Zero-based indexing is used. This encoding implicitly assumes column-major ordering of the elements. Secondly, rather than using a simple linked list or an array based representation, the list of the tuples is stored as a Red-Black Tree (RBT), a self-balancing binary search tree [7].

Briefly, an RBT is a collection of nodes, with each node containing the 2-tuple described above and links to two children nodes. There are two constraints: (i) each link points to a unique child node, and (ii) there are no links to the root node. The *index* within each 2-tuple is used as the key to identify each node. An example of this structure for a simple sparse matrix is shown in Figure 2(c). The ordering of the nodes and height of the tree (number of node levels below the root node) is controlled so that searching for a specific index (ie., retrieving an element at a specific location) has worst-case complexity of $\mathcal{O}(\log N)$. Insertion and removal of nodes (ie., matrix elements), also has the worst-case complexity of $\mathcal{O}(\log N)$. If a node to be inserted is known to have the largest index so far (eg., during incremental matrix construction), the search for where to place the node can be omitted, which in practice can considerably speed up the insertion process.

With the above element encoding, traversing an RBT in an ordered fashion (from the smallest to largest index) is equivalent to reading the elements in column-major ordering. This in turn allows for quick conversion of matrix data stored in RBT format into CSC format. The location of each element is simply decoded via $row = (index \bmod n_rows)$, and $column = \lfloor index/n_rows \rfloor$, where, for clarity, $\lfloor z \rfloor$ is the integer version of z , rounded towards zero. These operations are accomplished via direct integer arithmetic on CPUs. More details on the conversion are given in Section 5.

Within the hybrid storage framework, the RBT format is used for incremental construction of sparse matrices, either in an ordered or unordered fashion, and a subset of element-wise operations (such as inplace addition of values to specified elements). This in turn enables users to construct sparse matrices in the same way they might construct dense matrices—for instance, a loop over elements to be inserted without regard to storage format.

While the RBT format allows for fast element insertion, it is less suited than CSC for efficient linear algebra operations. The CSC format allows for exploitation of fast caches in modern CPUs due to the consecutive storage of non-zero elements in memory [24]. In contrast, accessing consecutive elements in the RBT format requires traversing the tree (following links from node to node), which in turn entails accessing node data that is not guaranteed to be consecutively stored in memory. Furthermore, obtaining the column and row indices requires explicit decoding of the index stored in each node, rather than a simple lookup in the CSC format.

4.3 Coordinate List Representation (COO)

The Coordinate List (COO) is a general concept where a list $L = (l_1, l_2, \dots, l_N)$ of 3-tuples represents the non-zero elements in a matrix. Each 3-tuple contains the location indices and value of the element, ie., $l = (\text{row}, \text{column}, \text{value})$. The format does not prescribe any ordering of the elements, and a simple linked list [7] can be used to represent L . However, in a computational implementation geared towards linear algebra operations [26], L is often represented as a set of three arrays:

1. The *values* array, which is a contiguous array of N floating point numbers holding the non-zero elements of the matrix.
2. The *rows* array, a contiguous array of N integers holding the row index of the corresponding value.
3. The *columns* array, a contiguous array of N integers holding the column index of the corresponding value.

As per the CSC format, all arrays use zero-based indexing, ie., the initial position in each array is 0. The elements in each array are sorted in column-major order for efficient lookup.

The array-based representation of COO is related to CSC, with the main difference that for each element the column indices are explicitly stored. This leads to the primary advantage of the COO format: it can greatly simplify the implementation of matrix processing algorithms. It also tends to be a natural format many non-expert users expect when first encountering sparse matrices. However, due to the explicit representation of column indices, the COO format contains redundancy and is hence less efficient (spacewise) than CSC for representing sparse matrices. An example of this is shown in Figure 2(d).

To contrast the differences in effort required in implementing matrix processing algorithms in CSC and COO, let us consider the problem of sparse matrix transposition. When using the COO format this is trivial to implement: simply swap the *rows* array with the *columns* array and then re-sort the elements so that column-major ordering is maintained. However, the same task for the CSC format is considerably more specialised: an efficient implementation in CSC would likely use an approach such as the elaborate `TRANSP` algorithm by Bank and Douglas [4], which is described through a 47-line pseudocode algorithm with annotations across two pages of text.

Our initial implementation of sparse matrix transposition used the COO based approach. COO was used simply due to shortage of available time for development and the need to flesh out other parts of sparse matrix functionality. When time allowed, we reimplemented sparse matrix transposition to use the abovementioned `TRANSP` algorithm. This resulted in considerable speedups, due to no longer requiring the time-consuming sort operation. We verified that the new CSC-based implementation is correct by comparing its output against the previous COO-based implementation on a large set of test matrices.

The relatively straightforward nature of COO format hence makes it well-suited for: (i) functionality contributed by time-constrained and/or non-expert users, (ii) relatively complex and/or less-common sparse matrix operations, and (iii) verifying the correct implementation of algorithms in the more complex CSC format. The volunteer driven nature of the Armadillo project makes its vibrancy and vitality depend in part on contributions received from users and the maintainability of the codebase. The number of core developers is small (ie., the authors of this paper), and hence difficult-to-understand or difficult-to-maintain code tends to be avoided, since the resources are simply not available to handle that burden.

The COO format is currently employed for less-commonly used tasks that involve bulk coordinate transformations, such as `reverse()` for flipping matrices column- or row-wise, and `repelem()`, where a matrix is generated by replicating each element several times from a given matrix. While it is certainly possible to adapt these functions to directly use the more complex CSC format, at the time of writing we have spent our time-constrained efforts on optimising and debugging more commonly used parts of the sparse matrix class.

5 Automatic Conversion Between Storage Formats

To circumvent the problems associated with selection and manual conversion between storage formats, our sparse matrix class employs a hybrid storage framework that *automatically* and *seamlessly* switches between the formats described in Section 4. By default, matrix elements are stored in CSC format. When needed, data in CSC format is internally converted to either the RBT or COO format, on which an operation or set of operations is performed. The matrix is automatically converted ('synced') back to the CSC format the next time an operation requiring the CSC format is performed.

<pre> X = scipy.sparse.rand(1000, 1000, 0.01) # manually convert to LIL format # to allow insertion of elements X = X.tolil() X[1,1] = 1.23 X[3,4] += 4.56 # random dense vector V = numpy.random.rand((1000)) # manually convert X to CSC format # for efficient multiplication X = X.tocsc() W = V * X </pre>	<pre> sp_mat X = sprandu(1000, 1000, 0.01); // automatic conversion to RBT format // for fast insertion of elements X(1,1) = 1.23; X(3,4) += 4.56; // random dense vector rowvec V(1000, fill::randu); // automatic conversion of X to CSC // prior to multiplication rowvec W = V * X; </pre>
--	---

Figure 3: Left panel: a Python program using the SciPy toolkit, requiring explicit conversions between sparse format types to achieve efficient execution; if an unsuitable sparse format is used for a given operation, SciPy will emit *TypeError* or *SparseEfficiencyWarning*. Right panel: A corresponding C++ program using the sparse matrix class, with the format conversions automatically done by the class.

The storage details and conversion operations are completely hidden from the user, who may not necessarily be knowledgeable about (or care to learn about) sparse matrix storage formats. This allows for simplified user code that focuses on high-level algorithm logic, which in turn increases readability and lowers maintenance. In contrast, other toolkits without automatic format conversion can cause either slow execution (as a non-optimal storage format might be used), or require many manual conversions. As an example, Figure 3 shows a short Python program using the SciPy toolkit [25] and a corresponding C++ program using the hybrid sparse matrix class. Manually initiated format conversions are required for efficient execution in the SciPy version; this causes both development time and code required to increase. If the user does not carefully consider the type of their sparse matrix at all times, they are likely to write inefficient code. In contrast, in the C++ program the format conversion is done automatically and behind the scenes.

A potential drawback of the automatic conversion between formats is the added computational cost. However, it turns out that COO/CSC conversions can be done in time that is linear in the number of non-zero elements in the matrix, and that CSC/RBT conversions can be done at worst in log-linear time. Since most sparse matrix operations are more expensive (eg., matrix multiplication), the conversion overhead turns out to be mostly negligible in practice. Below we present straightforward algorithms for conversion and note their asymptotic complexity in terms of the \mathcal{O} notation [7]. This is followed by discussing practical considerations that are not directly taken into account by the \mathcal{O} notation.

5.1 Conversion Between COO and CSC

Since the COO and CSC formats are quite similar, the conversion algorithms are straightforward. In fact the only parts of the formats to be converted are the *columns* and *col_offsets* arrays with the *rows* and *values* arrays remaining unchanged.

The algorithm for converting COO to CSC is given in Figure 4(a). In summary, the algorithm first determines the number of elements in each column (lines 6-8), and then ensures that the values in the *col_offsets* array are consecutively increasing (lines 9-10) so that they indicate the starting index of elements belonging to each column within the *values* array. The operations listed on line 5 and lines 9-10 each have a complexity of approximately $\mathcal{O}(n_{cols})$, while the operation listed on lines 6-8 has a complexity of $\mathcal{O}(N)$, where N is the number of non-zero elements in the matrix and n_{cols} is the number of columns. The complexity is hence $\mathcal{O}(N + 2n_{cols})$. As in most applications the number of non-zero elements will be considerably greater than the number of columns; the overall asymptotic complexity in these cases is $\mathcal{O}(N)$.

The corresponding algorithm for converting CSC to COO is shown in Figure 4(b). In essence the *col_offsets* array is unpacked into a *columns* array with length N . As such, the asymptotic complexity of this operation is $\mathcal{O}(N)$.

```

1 proc COO_to_CSC
2 input:  $N, n\_cols$  (integer scalars)
3 input:  $values, rows, columns$  (COO arrays)
4 allocate array  $col\_offsets$  with length  $n\_cols + 1$ 
5 forall  $j \in [0, n\_cols]$ :  $col\_offsets[j] \leftarrow 0$ 
6 forall  $i \in [0, N)$ :
7    $j \leftarrow columns[i] + 1$ 
8    $col\_offsets[j] \leftarrow col\_offsets[j] + 1$ 
9 forall  $j \in [1, n\_cols]$ :
10   $col\_offsets[j] \leftarrow col\_offsets[j] + col\_offsets[j-1]$ 
11 output:  $values, rows, col\_offsets$  (CSC arrays)

```

(a)

```

1 proc CSC_to_COO
2 input:  $N, n\_cols$  (integer scalars)
3 input:  $values, rows, col\_offsets$  (CSC arrays)
4 allocate array  $columns$  with length  $N$ 
5  $k \leftarrow 0$ 
6 forall  $j \in [0, n\_cols)$ :
7    $M \leftarrow col\_offsets[j+1] - col\_offsets[j]$ 
8   forall  $l \in [0, M)$ :
9      $columns[k+l] \leftarrow j$ 
10     $k \leftarrow k + M$ 
11 output:  $values, rows, columns$  (COO arrays)

```

(b)

Figure 4: Algorithms for: (a) conversion from COO to CSC, and (b) conversion from CSC to COO. Matrix elements in COO format are assumed to be stored in column-major ordering. All arrays and matrix locations use zero-based indexing. N indicates the number of non-zero elements, while n_cols indicates the number of columns. Details for the CSC and COO arrays are given in Section 4.

```

1 proc CSC_to_RBT
2 input:  $N, n\_rows, n\_cols$  (integer scalars)
3 input:  $values, rows, col\_offsets$  (CSC arrays)
4 declare red-black tree  $T$ 
5 forall  $j \in [0, n\_cols)$ :
6    $start \leftarrow col\_offsets[j]$ 
7    $end \leftarrow col\_offsets[j+1]$ 
8   forall  $k \in [start, end)$ :
9      $index \leftarrow row\_indices[k] + j * n\_rows$ 
10     $l \leftarrow (index, values[k])$ 
11    insert node  $l$  into  $T$ 
12 output:  $T$  (red-black tree)

```

(a)

```

1 proc RBT_to_CSC
2 input:  $N, n\_rows, n\_cols$  (integer scalars)
3 input:  $T$  (red-black tree)
4 allocate array  $values$  with length  $N$ 
5 allocate array  $row\_indices$  with length  $N$ 
6 allocate array  $col\_offsets$  with length  $n\_cols + 1$ 
7 forall  $j \in [0, n\_cols)$ :  $col\_offsets[j] \leftarrow 0$ 
8  $k \leftarrow 0$ 
9 foreach node  $l \in T$ , where  $l = (index, value)$ :
10   $values[k] \leftarrow value$ 
11   $row\_indices[k] \leftarrow index \bmod n\_rows$ 
12   $j \leftarrow \lfloor index / n\_rows \rfloor$ 
13   $col\_offsets[j+1] \leftarrow col\_offsets[j+1] + 1$ 
14   $k \leftarrow k + 1$ 
15 forall  $j \in [1, n\_cols]$ :
16   $col\_offsets[j] \leftarrow col\_offsets[j] + col\_offsets[j-1]$ 
17 output:  $values, rows, col\_offsets$  (CSC arrays)

```

(b)

Figure 5: Algorithms for: (a) conversion from CSC to RBT, and (b) conversion from RBT to CSC. All arrays and matrix locations use zero-based indexing. N indicates the number of non-zero elements, while n_rows and n_cols indicate the number of rows and columns, respectively. Details for the CSC arrays are given in Section 4.

5.2 Conversion Between CSC and RBT

The conversion between the CSC and RBT formats is also straightforward and can be accomplished using the algorithms shown in Figure 5. In essence, the CSC to RBT conversion involves encoding the location of each matrix element to a linear index, followed by inserting a node with that index and the corresponding element value into the RBT. The worst-case complexity for inserting all elements into an RBT is $\mathcal{O}(N \cdot \log N)$. However, as the elements in the CSC format are guaranteed to be stored according to column-major ordering (as per Section 4.1), and the location encoding assumes column-major ordering (as per Section 4.2), the insertion of a node into an RBT can be accomplished without searching for the node location. While the worst-case cost of $\mathcal{O}(N \cdot \log N)$ is maintained due to tree maintenance (ie., controlling the height of the tree) [7], in practice the amortised insertion cost is typically lower due to avoidance of the search.

Converting an RBT to CSC involves traversing through the nodes of the tree from the lowest to highest index, which is equivalent to reading the elements in column-major format. The value stored in each node is hence simply copied into the corresponding location in the CSC $values$ array. The index stored in each node is decoded into row and column indices, as per Section 4.2, with the CSC $row_indices$ and $col_offsets$ arrays adjusted accordingly. The worst-case cost for finding each element in the RBT is $\mathcal{O}(\log N)$, which results in the asymptotic worst-case cost of $\mathcal{O}(N \cdot \log N)$ for the whole conversion. However, in practice most

consecutive elements are in nearby nodes, which on average reduces the number of traversals across nodes, resulting in considerably lower amortised conversion cost.

5.3 Practical Considerations

Since the conversion algorithms given in Figures 4 and 5 are quite straightforward, the \mathcal{O} notation does not hide any large constant factors. For COO/CSC conversions the cost is $\mathcal{O}(N)$, while for CSC/RBT conversions the worst-case cost is $\mathcal{O}(N \cdot \log N)$. In contrast, many mathematical operations on sparse matrices have much higher computational cost than the conversion algorithms. Even simply adding two sparse matrices can be much more expensive than a conversion. Although the addition operation still takes $\mathcal{O}(N)$ time (assuming N is identical for both matrices), there is a lot of hidden constant overhead, since the sparsity pattern of the resulting matrix must be computed first [26]. A similar situation applies for multiplication of two sparse matrices, whose the cost is in general superlinear in the number of non-zeros of either input matrix [4, 9]. Sparse matrix factorisations, such as LU factorisation [16, 20] and Cholesky factorisation [15], are also much more expensive than the conversion overhead. Other factorisations and higher-level operations can exhibit similar complexity characteristics. Given this, the cost of format conversions is heavily outweighed by the user convenience that they allow.

6 Empirical Evaluation

To demonstrate the advantages of the hybrid storage framework and the template-based expression optimisation mechanism, we have performed a set of experiments, measuring the wall-clock time (elapsed real time) required for:

1. Unordered element insertion into a sparse matrix, where the elements are inserted at random locations in random order.
2. Quasi-ordered element insertion into a sparse matrix, where each new inserted element is at a random location that is past the previously inserted element, under the constraint of column-major ordering.
3. Calculation of $\text{trace}(A^T B)$, where A and B are randomly generated sparse matrices.
4. Obtaining a diagonal matrix from the $(A + B)$ expression, where A and B are randomly generated sparse matrices.

In all cases the sparse matrices have a size of $10,000 \times 10,000$, with four settings for the density of non-zero elements: 0.01%, 0.1%, 1%, 10%. The experiments were done on a machine with an Intel Xeon E5-2630L CPU running at 2 GHz, using the GCC v5.4 compiler. Each experiment was repeated 10 times, and the average wall-clock time is reported. The wall-clock time measures the total time taken from the start to the end of each run, and includes necessary overheads such as memory allocation.

Figure 6 shows the average wall-clock time taken for element insertion done directly using the underlying storage formats (ie., CSC, COO, RBT, as per Section 4), as well as the hybrid approach which uses RBT followed by conversion to CSC. The CSC and COO formats use oversized storage as a form of optimisation (as mentioned in Section 4.1), where the underlying arrays are grown in chunks of 1024 elements in order to reduce both the number of memory reallocations and array copy operations due to element insertions.

In all cases bar one, the RBT format is the quickest for insertion, generally by one or two orders of magnitude. The conversion from RBT to CSC adds negligible overhead. For the single case of quasi-ordered insertion to reach the density of 0.01%, the COO format is slightly quicker than RBT. This is due to the relatively simple nature of the COO format, as well as the ordered nature of the element insertion where the elements are directly placed into the oversized COO arrays (ie., no sorting required). Furthermore, due to the very low density of non-zero elements and the chunked nature of COO array growth, the number of reallocations of the COO arrays is relatively low. In contrast, inserting a new element into RBT requires the allocation of memory for a new node, and modifying the tree to append the node. For larger densities ($\geq 0.1\%$), the COO element insertion process quickly becomes more time consuming than RBT element insertion, due to an increased amount of array reallocations and the increased size of the copied arrays. Compared to COO, the CSC format is more complex and has the additional burden of recalculating the column offsets (*col_offsets*) array for each inserted element.

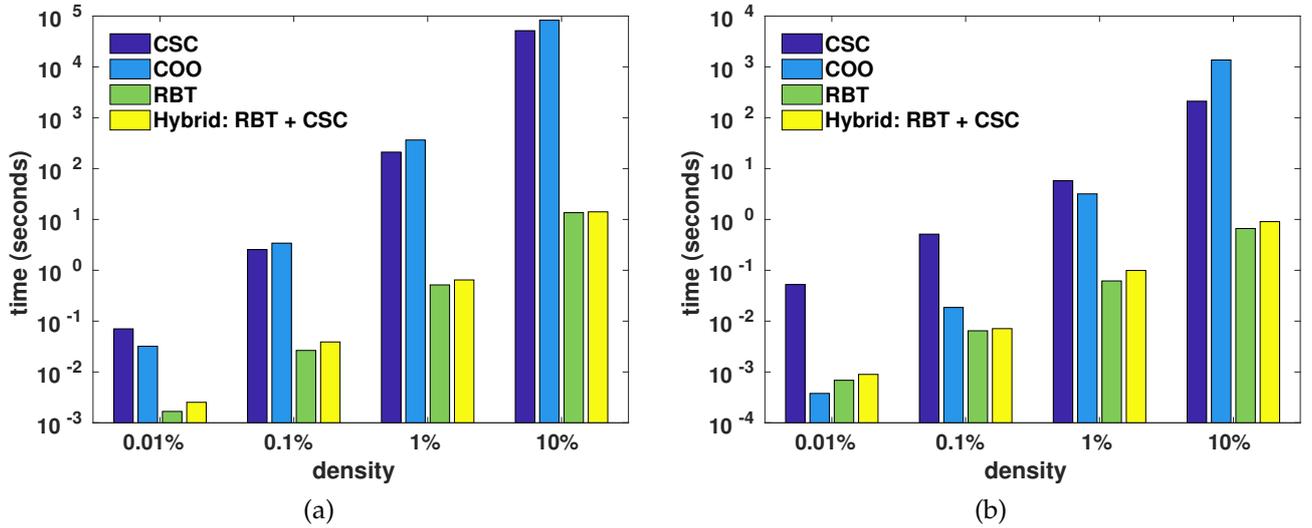


Figure 6: Wall-clock time taken to insert elements into a $10,000 \times 10,000$ sparse matrix to achieve various densities of non-zero elements. In (a), the elements are inserted at random locations in random order. In (b), the elements are inserted in a quasi-ordered fashion, where each new inserted element is at a random location that is past the previously inserted element, using column-major ordering.

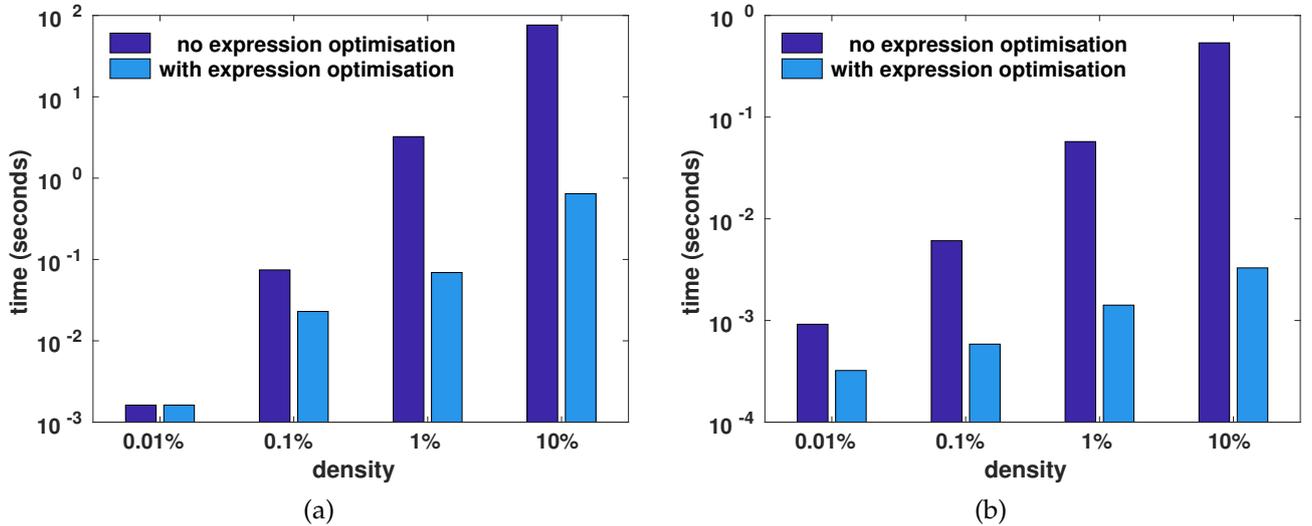


Figure 7: Wall-clock time taken to calculate the expressions (a) $\text{trace}(A.t()*B)$ and (b) $\text{diagmat}(A+B)$, where A and B are randomly generated sparse matrices with a size of $10,000 \times 10,000$ and various densities of non-zero elements. The expressions were calculated with and without the aid of the template-based optimisation of compound expression described in Section 3. As per Table 1, $x.t()$ returns the transpose of matrix x , while $\text{diagmat}(x)$ returns a diagonal matrix constructed from the main diagonal of x .

Figure 7 shows the wall-clock time taken to calculate the expressions $\text{trace}(A.t()*B)$ and $\text{diagmat}(A+B)$, with and without the aid of the automatic template-based optimisation of compound expression described in Section 3. For both expressions, employing expression optimisation leads to considerable reduction in the wall-clock time. As the density increases (ie., more non-zero elements), more time is saved via expression optimisation.

For the $\text{trace}(A.t()*B)$ expression, the expression optimisation computes the trace by omitting the explicit transpose operation and performing a partial matrix multiplication to obtain only the diagonal elements. In a similar fashion, the expression optimisation for the $\text{diagmat}(A+B)$ expression directly generates the diagonal matrix by performing a partial matrix addition, where only the diagonal elements of the two matrices are added. As well as avoiding full matrix addition, the generation of a temporary intermediary matrix to hold the complete result of the matrix addition is also avoided.

7 Conclusion

Driven by a scarcity of easy-to-use tools for algorithm development that requires use of sparse matrices, we have devised a practical sparse matrix class for the C++ language. The sparse matrix class internally uses a hybrid storage framework, which automatically and seamlessly switches between several underlying formats, depending on which format is best suited and/or available for specific operations. This allows the user to write sparse linear algebra without requiring to consider the intricacies and limitations of various storage formats. Furthermore, the sparse matrix class employs a template meta-programming framework that can automatically optimise several common expression patterns, resulting in faster execution.

The source code for the sparse matrix class and its associated functions is included in recent releases of the cross-platform and open-source Armadillo linear algebra library [27], available from <http://arma.sourceforge.net>. The code is provided under the permissive Apache 2.0 license [30], allowing unencumbered use in both open-source and proprietary projects (eg. product development).

The sparse matrix class has already been successfully used in open-source projects such as the *mlpack* library for machine learning [8], and the *ensmallen* library for mathematical function optimisation [5]. In both cases the sparse matrix class is used to allow various algorithms to be run on either sparse or dense datasets. Furthermore, bi-directional bindings for the class are provided to the R environment via the Rcpp bridge [13]. Avenues for further exploration include expanding the hybrid storage framework with more sparse matrix formats [26, 11, 3] in order to provide speedups for specialised use cases.

Acknowledgements

We would like to thank the anonymous reviewers, as well as colleagues at the University of Queensland (Ian Hayes, George Havas, Arnold Wiliem) and Data61/CSIRO (Dan Pagendam, Josh Bowden, Regis Riveret) for discussions leading to the improvement of this article.

References

- [1] D. Abrahams and A. Gurtovoy. *C++ Template Metaprogramming: Concepts, Tools, and Techniques from Boost and Beyond*. Addison-Wesley Professional, 2004.
- [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] Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst. *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*. SIAM, 2000.
- [4] R. E. Bank and C. C. Douglas. Sparse matrix multiplication package (SMMP). *Advances in Computational Mathematics*, 1(1):127–137, 1993.
- [5] S. Bhardwaj, R. Curtin, M. Edel, Y. Mentekidis, and C. Sanderson. *ensmallen*: a flexible C++ library for efficient function optimization. *Workshop on Systems for ML and Open Source Software at NIPS / NeurIPS*, 2018.
- [6] N. Boumal, V. Voroninski, and A. Bandeira. The non-convex Burer-Monteiro approach works on smooth semidefinite programs. In *Advances in Neural Information Processing Systems*, pages 2757–2765, 2016.
- [7] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 3rd edition, 2009.
- [8] R. Curtin, M. Edel, M. Lozhnikov, Y. Mentekidis, S. Ghaisas, and S. Zhang. *mlpack 3*: a fast, flexible machine learning library. *Journal of Open Source Software*, 3:726, 2018.
- [9] T. A. Davis. *Direct methods for sparse linear systems*. SIAM, 2006.
- [10] T. A. Davis, S. Rajamanickam, and W. M. Sid-Lakhdar. A survey of direct methods for sparse linear systems. *Acta Numerica*, 25:383–566, 2016.
- [11] I. S. Duff, A. M. Erisman, and J. K. Reid. *Direct methods for sparse matrices*. Oxford University Press, 2nd edition, 2017.
- [12] J. W. Eaton, D. Bateman, S. Hauberg, and R. Wehbring. *GNU Octave 4.2 Reference Manual*. Samurai Media Limited, 2017.
- [13] D. Eddelbuettel and C. Sanderson. RcppArmadillo: Accelerating R with high-performance C++ linear algebra. *Computational Statistics & Data Analysis*, 71:1054–1063, 2014.
- [14] L. El Ghaoui and H. Le Bret. Robust solutions to least-squares problems with uncertain data. *SIAM Journal on Matrix Analysis and Applications*, 18(4):1035–1064, 1997.
- [15] A. George and E. Ng. On the complexity of sparse QR and LU factorization of finite-element matrices. *SIAM Journal on Scientific and Statistical Computing*, 9(5):849–861, 1988.
- [16] J. R. Gilbert, X. S. Li, E. G. Ng, and B. W. Peyton. Computing row and column counts for sparse QR and LU factorization. *BIT Numerical Mathematics*, 41(4):693–710, 2001.

- [17] K. Iglberger, G. Hager, J. Treibig, and U. Rde. Expression templates revisited: a performance analysis of current methodologies. *SIAM Journal on Scientific Computing*, 34(2):C42–C69, 2012.
- [18] G. R. Lanckriet, N. Cristianini, P. Bartlett, L. E. Ghaoui, and M. I. Jordan. Learning the kernel matrix with semidefinite programming. *Journal of Machine Learning Research*, 5:27–72, 2004.
- [19] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*. SIAM, 1998.
- [20] X. S. Li. An overview of SuperLU: Algorithms, implementation, and user interface. *ACM Transactions on Mathematical Software (TOMS)*, 31(3):302–325, 2005.
- [21] P. Liniker, O. Beckmann, and P. H. Kelly. Delayed evaluation, self-optimising software components as a programming model. In *European Conference on Parallel Processing - Euro-Par 2002. Lecture Notes in Computer Science (LNCS)*, volume 2400, pages 666–673, 2002.
- [22] MathWorks. MATLAB Documentation - Accessing Sparse Matrices. <https://www.mathworks.com/help/matlab/math/accessing-sparse-matrices.html>, 2018.
- [23] M. Mernik, J. Heering, and A. M. Sloane. When and how to develop domain-specific languages. *ACM Computing Surveys*, 37(4):316–344, 2005.
- [24] S. Mittal. A survey of recent prefetching techniques for processor caches. *ACM Computing Surveys*, 49(2):35:1–35:35, 2016.
- [25] J. Nunez-Iglesias, S. van der Walt, and H. Dashnow. *Elegant SciPy: The Art of Scientific Python*. O'Reilly Media, 2017.
- [26] Y. Saad. SPARSKIT: A basic tool kit for sparse matrix computations. Technical Report NASA-CR-185876, NASA Ames Research Center, 1990.
- [27] C. Sanderson and R. Curtin. Armadillo: a template-based C++ library for linear algebra. *Journal of Open Source Software*, 1:26, 2016.
- [28] C. Sanderson and R. Curtin. A user-friendly hybrid sparse matrix class in C++. In *Mathematical Software - ICMS 2018. Lecture Notes in Computer Science (LNCS)*, volume 10931, pages 422–430, 2018.
- [29] M. Scherr and S. Chiba. Almost first-class language embedding: taming staged embedded DSLs. In *ACM SIGPLAN International Conference on Generative Programming: Concepts and Experiences*, pages 21–30, 2015.
- [30] A. St. Laurent. *Understanding Open Source and Free Software Licensing*. O'Reilly Media, 2008.
- [31] B. Stroustrup. *The C++ Programming Language*. Addison-Wesley, 4th edition, 2013.
- [32] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95, 1996.
- [33] D. Vandevorde and N. M. Josuttis. *C++ Templates: The Complete Guide*. Addison-Wesley, 2nd edition, 2017.
- [34] T. L. Veldhuizen. C++ templates as partial evaluation. In *ACM SIGPLAN Workshop on Partial Evaluation and Semantics-Based Program Manipulation*, pages 13–18, 1999.