

Article

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

Conrad Sanderson ^{1,2,4,*} and Ryan Curtin ^{3,4}

¹ Data61 / CSIRO, Australia

² University of Queensland, Australia

³ RelationalAI, USA

⁴ Arroyo Consortium

* Correspondence: conradsand@ieee.org

† This paper is an extended version of our paper published in International Congress on Mathematical Software 2018.

Version 6th July 2019 submitted to Math. Comput. Appl.

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

MSC: 68N99; 65Y04; 65Y15; 65F50

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 [1] 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,

30 with severe runtime penalties if it is not used in the right way. This can be challenging and bewildering
31 for users who simply want to create and use sparse matrices, and do not have the time, expertise, or
32 desire to understand the advantages and disadvantages of each format. To achieve good performance,
33 several formats may need to be used in one program, requiring explicit selection and conversion
34 between the formats. This multitude of sparse matrix classes complicates the programming task, adds
35 to the maintenance burden, and increases the likelihood of bugs.

36 Driven by the above concerns, we have devised a practical and user-friendly sparse matrix class
37 for the C++ language [2]. The sparse matrix class uses a hybrid storage framework, which *automatically*
38 and *seamlessly* switches between three data storage formats, depending on which format is best suited
39 and/or available for specific operations:

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

47 The COO format is important to point out, as the source code for the sparse matrix class is
48 distributed and maintained as part of the open-source Armadillo library [3]. Due to its simpler nature,
49 the COO format facilitates functionality contributions from time-constrained and/or non-expert users,
50 as well as reducing maintenance and debugging overhead for the library maintainers.

51 While there are many other sparse matrix implementations in existence, to our knowledge the
52 presented approach is the first to offer a unified interface with automatic format switching under the
53 hood. Most toolkits are limited to either a single format or multiple formats the user must manually
54 convert between. The comprehensive SPARSKIT package [4] contains 16, and SciPy contains seven
55 formats [1]. In these toolkits the user must manually convert between formats. On the other hand,
56 both MATLAB and GNU Octave [5] contain sparse matrix implementations, but they supply only the
57 CSC format [6], meaning that users must write their code in special ways to ensure its efficiency [7].
58 This is a similar situation to the Blaze library (bitbucket.org/blaze-lib/blaze) [8], which implements
59 only a CSR/CSC format sparse matrix. Users are explicitly discouraged from individual element
60 insertions and, for efficiency, must construct their sparse matrices in the restricted environment of batch
61 insertion. The Eigen C++ matrix library (eigen.tuxfamily.org) uses a specialised sparse matrix format
62 which has deliberate redundancy and overprovisioned storage. While this can help with reducing
63 the computational effort of element insertion in some situations, it requires manual care to maintain
64 storage efficiency. Furthermore, as the cost of random insertion of elements is still high, the associated
65 documentation recommends to manually construct a COO-like representation of all the elements,
66 from which the actual sparse matrix is then constructed. The IT++ library (itpp.sourceforge.net)
67 has a cumbersome sparse matrix class with a custom format that also employs overprovisioned
68 storage. The format is less efficient storage-wise than CSC unless explicit manual care is taken. Data
69 is stored in unordered fashion which allows for faster element insertion than CSC, but at the cost
70 of reduced performance for linear algebra operations. Thus, overall, the landscape of sparse matrix
71 implementations is composed of libraries where a user must be aware of some of the internal storage
72 details of these implementations in order to produce efficient code; this is not ideal.

73 To make the situation even more complex, there are also numerous other sparse matrix formats [4,
74 9]. Examples are the modified compressed row/column format (intended for sparse matrices with all
75 non-zero elements on the diagonal), block compressed storage format (intended for sparse matrices
76 with dense submatrices), compressed diagonal format (intended for straightforward storage of banded
77 sparse matrices under the assumption of constant bandwidth), and the skyline format (intended for
78 more efficient storage of banded sparse matrices with irregular bandwidth). As these formats are

79 focused on specialised use cases, their utility is typically not very general. Thus we have currently
 80 opted against including these formats in our hybrid framework, though it would be relatively easy to
 81 accommodate more formats in the future.

82 To further promote efficient execution, the sparse matrix class internally implements a delayed
 83 evaluation framework [10] based on template meta-programming [11,12] combined with operator
 84 overloading [2]. In delayed evaluation, the evaluation of a given compound mathematical expression
 85 is delayed until its value is required (ie., assigned to a variable). This is in contrast to eager evaluation
 86 (also known as strict evaluation), where each component of a compound expression is evaluated
 87 immediately. As such, the delayed evaluation framework allows automatic compile-time analysis of
 88 compound expressions, which in turns allows for automatic detection and optimisation of common
 89 expression patterns. For example, several operations can be combined to reduce the required
 90 computational effort.

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

96 The paper is continued as follows. In Section 2 we overview the functionality provided by the
 97 sparse matrix class and its associated functions. The delayed evaluation approach is overviewed
 98 in Section 3. In Section 4 we describe the underlying storage formats used by the class, and the
 99 scenarios that each of the formats is best suited for. In Section 5 we discuss the costs for switching
 100 between the formats. Section 6 provides an empirical evaluation showing the advantages of the hybrid
 101 storage framework and the delayed evaluation approach. The salient points and avenues for further
 102 exploitation are summarised in Section 7. This article is a thoroughly revised and extended version of
 103 our earlier work [13].

104 2. Functionality

105 The sparse matrix class and its associated functions provide a user-friendly suite of essential
 106 sparse linear algebra functionality, including fundamental operations such as addition, matrix
 107 multiplication and submatrix manipulation. The class supports storing elements as integers,
 108 single- and double-precision floating point numbers, as well as complex numbers. Various sparse
 109 eigendecompositions and linear equation solvers are provided through integration with low-level
 110 routines in the de-facto standard ARPACK [14] and SuperLU libraries [15]. The resultant high-level
 111 functions automatically take care of tedious and cumbersome details such as memory management,
 112 allowing the user to concentrate their programming effort on mathematical details.

113 C++ language features such as overloading of operators (eg., * and +) [2] are exploited to allow
 114 mathematical operations with matrices to be expressed in a concise and easy-to-read manner, in a
 115 similar fashion to the proprietary MATLAB language. For example, given sparse matrices A, B, and C,
 116 a mathematical expression such as

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

117 can be written directly in C++ as

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

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

120 The aggregate of the sparse matrix class, operator overloading and associated functions on sparse
 121 matrices is an instance of a Domain Specific Language (sparse linear algebra in this case) embedded
 122 within the host C++ language [16,17]. This allows complex algorithms relying on sparse matrices to be
 123 easily developed and integrated within a larger C++ program, making the sparse matrix class directly
 124 useful in application/product development.

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 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.sf.net/docs.html#SpMat> for more detailed documentation.

```
#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`).

125 3. Template-Based Optimisation of Compound Expressions

126 The sparse matrix class uses a delayed evaluation approach, allowing several operations to be
 127 combined to reduce the amount of computation and/or temporary objects. In contrast to brute-force
 128 evaluations, delayed evaluation can provide considerable performance improvements as well as
 129 reduced memory usage [18]. The delayed evaluation machinery is accomplished through template
 130 meta-programming [11,12], where a type-based signature of a compound expression (set of consecutive
 131 mathematical operations) is automatically constructed. The C++ compiler is then automatically
 132 induced to detect common expression patterns at compile time, followed by selecting the most
 133 computationally efficient implementations.

134 As an example of the possible efficiency gains, let us consider the expression $\text{trace}(A.t() * B)$,
 135 which often appears as a fundamental quantity in semidefinite programs [19]. These computations are
 136 thus used in a wide variety of diverse fields, most notably machine learning [20–22]. A brute-force
 137 implementation would evaluate the transpose first, $A.t()$, and store the result in a temporary matrix
 138 $T1$. The next operation would be a time consuming matrix multiplication, $T1 * B$, with the result
 139 stored in another temporary matrix $T2$. The trace operation (sum of diagonal elements) would then be
 140 applied on $T2$. The explicit transpose, full matrix multiplication and creation of the temporary matrices
 141 is suboptimal from an efficiency point of view, as for the trace operation we require only the diagonal
 142 elements of the $A.t() * B$ expression.

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

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

$$Op<sp_mat, op_trans>$$

152 where $Op<...>$ indicates that Op is a template class, with the items between ' $<$ ' and ' $>$ ' specifying
 153 template parameters. In this case the $Op<sp_mat, op_trans>$ object type indicates that a reference
 154 to a matrix is stored and that a transpose operation is requested. In turn, the compound expression
 155 $A.t() * B$ is converted to an instance of the lightweight $Glue$ object with the following type:

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

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

160 There are several implemented forms of the $\text{trace}()$ function, one of which is automatically
 161 chosen by the C++ compiler to handle the $Glue< Op<sp_mat, op_trans>, sp_mat, glue_times>$
 162 expression. The specific form of $\text{trace}()$ takes references to the A and B matrices, and executes a *partial*
 163 matrix multiplication to obtain only the diagonal elements of the $A.t() * B$ expression. All of this
 164 is accomplished without generating temporary matrices. Furthermore, as the $Glue$ and Op objects
 165 only hold references, they are in effect optimised away by modern C++ compilers [12]: the resultant
 166 machine code appears as if the $Glue$ and Op objects never existed in the first place.

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




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 [2], we use zero-based indexing.

170 4. Storage Formats for Sparse Data

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

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

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

186 4.1. Compressed Sparse Column (CSC)

187 The CSC format [4] uses column-major ordering where the elements are stored column-by-column,
 188 with consecutive non-zero elements in each column stored consecutively in memory. Three arrays are
 189 used to represent a sparse matrix:

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

196 Following C++ convention [2], all arrays use zero-based indexing, ie., the initial position in each array
 197 is denoted by 0. For consistency, element locations within a matrix are also encoded as starting at zero,
 198 ie., the initial row and column are both denoted by 0. Furthermore, the row indices for elements in
 199 each column are kept sorted in ascending manner. In many applications, sparse matrices have more
 200 non-zero elements than the number of columns, leading to the *col_offsets* array being typically much
 201 smaller than the *values* array.

202 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
 203 the n -th entry in the *values* array as $v[n]$. The number of non-zero elements in column i is determined

204 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
205 non-zero elements, then the first element is obtained via $v[c[i]]$, and $r[c[i]]$ is the corresponding row of
206 the element. An example of this format is shown in Figure 2(b).

207 The CSC format is well-suited for efficient sparse linear algebra operations such as vector-matrix
208 multiplication. This is due to consecutive non-zero elements in each column being stored next to
209 each other in memory, which allows modern CPUs to speculatively read ahead elements from the
210 main memory into fast cache memory [23]. The CSC format is also suited for operations that do not
211 change the structure of the matrix, such as element-wise operations on the non-zero elements (eg.,
212 multiplication by a scalar). The format also affords relatively efficient random element access; to locate
213 an element (or determine that it is not stored), a single lookup to the beginning of the desired column
214 can be performed, followed by a binary search [24] through the *rows* array to find the element.

215 While the CSC format provides a compact representation yielding efficient execution of linear
216 algebra operations, it has two main disadvantages. The first disadvantage is that the design and
217 implementation of efficient algorithms for many sparse matrix operations (such as matrix-matrix
218 multiplication) tend to be non-trivial [4,25]. This stems not only from the sparse nature of the data, but
219 also due to the need to (i) explicitly keep track of the column offsets, and (ii) keep the row indices for
220 elements in each column sorted in ascending manner. In our experience, the process of designing and
221 implementing efficient matrix processing algorithms in CSC is a time-consuming affair — it is both
222 finicky and prone to subtle bugs.

223 The second disadvantage of CSC is the computational effort required to insert a new element [6].
224 In the worst-case scenario, memory for three new larger-sized arrays (containing the values and
225 locations) must first be allocated, the position of the new element determined within the arrays, data
226 from the old arrays copied to the new arrays, data for the new element placed in the new arrays, and
227 finally the memory used by the old arrays deallocated. As the number of elements in the matrix grows,
228 the entire process becomes slower.

229 There are opportunities for some optimisation, such as using oversized storage to reduce memory
230 allocations, where a new element past all the previous elements can be readily inserted. However, this
231 does not help when a new non-zero element is inserted between two existing non-zero elements. It
232 is also possible to perform batch insertions with some speedup by first sorting all the elements to be
233 inserted and then merging with the existing data arrays. While the above approaches can be effective,
234 they require the user to explicitly deal with cumbersome low-level storage details instead of focusing
235 on high-level functionality.

236 The CSC format was chosen over the related Compressed Sparse Row (CSR) format [4] for two
237 main reasons: (i) to ensure compatibility with external libraries such as the SuperLU solver [15], and
238 (ii) to ensure consistency with the surrounding infrastructure provided by the Armadillo library, which
239 uses column-major dense matrix representation to take advantage of low-level functions provided by
240 LAPACK [26].

241 4.2. Red-Black Tree (RBT)

242 To address the efficiency problems with element insertion at arbitrary locations, we first represent
243 each element as a 2-tuple, $l = (index, value)$, where *index* encodes the location of the element as
244 $index = row + column \times n_rows$. Zero-based indexing is used. This encoding implicitly assumes
245 column-major ordering of the elements. Secondly, rather than using a simple linked list or an array
246 based representation, the list of the tuples is stored as a Red-Black Tree (RBT), a self-balancing binary
247 search tree [24].

248 Briefly, an RBT is a collection of nodes, with each node containing the 2-tuple described above
249 and links to two children nodes. There are two constraints: (i) each link points to a unique child
250 node, and (ii) there are no links to the root node. The *index* within each 2-tuple is used as the key to
251 identify each node. An example of this structure for a simple sparse matrix is shown in Figure 2(c). The
252 ordering of the nodes and height of the tree (number of node levels below the root node) is controlled

253 so that searching for a specific index (ie., retrieving an element at a specific location) has worst-case
 254 complexity of $\mathcal{O}(\log N)$. Insertion and removal of nodes (ie., matrix elements), also has the worst-case
 255 complexity of $\mathcal{O}(\log N)$. If a node to be inserted is known to have the largest index so far (eg., during
 256 incremental matrix construction), the search for where to place the node can be omitted, which in
 257 practice can considerably speed up the insertion process.

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

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

269 While the RBT format allows for fast element insertion, it is less suited than CSC for efficient linear
 270 algebra operations. The CSC format allows for exploitation of fast caches in modern CPUs due to the
 271 consecutive storage of non-zero elements in memory [23]. In contrast, accessing consecutive elements
 272 in the RBT format requires traversing the tree (following links from node to node), which in turn
 273 entails accessing node data that is not guaranteed to be consecutively stored in memory. Furthermore,
 274 obtaining the column and row indices requires explicit decoding of the index stored in each node,
 275 rather than a simple lookup in the CSC format.

276 4.3. Coordinate List Representation (COO)

277 The Coordinate List (COO) is a general concept where a list $L = (l_1, l_2, \dots, l_N)$ of 3-tuples
 278 represents the non-zero elements in a matrix. Each 3-tuple contains the location indices and value of
 279 the element, ie., $l = (row, column, value)$. The format does not prescribe any ordering of the elements,
 280 and a simple linked list [24] can be used to represent L . However, in a computational implementation
 281 geared towards linear algebra operations [4], L is often represented as a set of three arrays:

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

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

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

297 To contrast the differences in effort required in implementing matrix processing algorithms in
 298 CSC and COO, let us consider the problem of sparse matrix transposition. When using the COO
 299 format this is trivial to implement: simply swap the *rows* array with the *columns* array and then re-sort
 300 the elements so that column-major ordering is maintained. However, the same task for the CSC format

301 is considerably more specialised: an efficient implementation in CSC would likely use an approach
302 such as the elaborate TRANSP algorithm by Bank and Douglas [25], which is described through a 47-line
303 pseudocode algorithm with annotations across two pages of text.

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

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

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

325 5. Automatic Conversion Between Storage Formats

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

332 The storage details and conversion operations are completely hidden from the user, who may
333 not necessarily be knowledgeable about (or care to learn about) sparse matrix storage formats. This
334 allows for simplified user code that focuses on high-level algorithm logic, which in turn increases
335 readability and lowers maintenance. In contrast, other toolkits without automatic format conversion
336 can cause either slow execution (as a non-optimal storage format might be used), or require many
337 manual conversions. As an example, Figure 3 shows a short Python program using the SciPy toolkit [1]
338 and a corresponding C++ program using the hybrid sparse matrix class. Manually initiated format
339 conversions are required for efficient execution in the SciPy version; this causes both development
340 time and code required to increase. If the user does not carefully consider the type of their sparse
341 matrix at all times, they are likely to write inefficient code. In contrast, in the C++ program the format
342 conversion is done automatically and behind the scenes.

343 A potential drawback of the automatic conversion between formats is the added computational
344 cost. However, it turns out that COO/CSC conversions can be done in time that is linear in the
345 number of non-zero elements in the matrix, and that CSC/RBT conversions can be done at worst in
346 log-linear time. Since most sparse matrix operations are more expensive (eg., matrix multiplication), the
347 conversion overhead turns out to be mostly negligible in practice. Below we present straightforward
348 algorithms for conversion and note their asymptotic complexity in terms of the \mathcal{O} notation [24]. This

<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.

349 is followed by discussing practical considerations that are not directly taken into account by the
350 \mathcal{O} notation.

351 5.1. Conversion Between COO and CSC

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

355 The algorithm for converting COO to CSC is given in Figure 4(a). In summary, the algorithm first
356 determines the number of elements in each column (lines 6-8), and then ensures that the values in
357 the *col_offsets* array are consecutively increasing (lines 9-10) so that they indicate the starting index of
358 elements belonging to each column within the *values* array. The operations listed on line 5 and lines
359 9-10 each have a complexity of approximately $\mathcal{O}(n_{\text{cols}})$, while the operation listed on lines 6-8 has
360 a complexity of $\mathcal{O}(N)$, where N is the number of non-zero elements in the matrix and n_{cols} is the
361 number of columns. The complexity is hence $\mathcal{O}(N + 2n_{\text{cols}})$. As in most applications the number of
362 non-zero elements will be considerably greater than the number of columns, the overall asymptotic
363 complexity in these cases is $\mathcal{O}(N)$.

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

367 5.2. Conversion Between CSC and RBT

368 The conversion between the CSC and RBT formats is also straightforward and can be
369 accomplished using the algorithms shown in Figure 5. In essence, the CSC to RBT conversion involves
370 encoding the location of each matrix element to a linear index, followed by inserting a node with that
371 index and the corresponding element value into the RBT. The worst-case complexity for inserting all
372 elements into an RBT is $\mathcal{O}(N \cdot \log N)$. However, as the elements in the CSC format are guaranteed to
373 be stored according to column-major ordering (as per Section 4.1), and the location encoding assumes
374 column-major ordering (as per Section 4.2), the insertion of a node into an RBT can be accomplished
375 without searching for the node location. While the worst-case cost of $\mathcal{O}(N \cdot \log N)$ is maintained due
376 to tree maintenance (ie., controlling the height of the tree) [24], in practice the amortised insertion cost
377 is typically lower due to avoidance of the search.

<pre> 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) </pre>	<pre> 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) </pre>
(a)	(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.

<pre> 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) </pre>	<pre> 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) </pre>
(a)	(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 row and columns, respectively. Details for the CSC arrays are given in Section 4.

378 Converting an RBT to CSC involves traversing through the nodes of the tree from the lowest to
379 highest index, which is equivalent to reading the elements in column-major format. The value stored
380 in each node is hence simply copied into the corresponding location in the CSC $values$ array. The
381 index stored in each node is decoded into row and column indices, as per Section 4.2, with the CSC
382 $row_indices$ and $col_offsets$ arrays adjusted accordingly. The worst-case cost for finding each element in
383 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
384 conversion. However, in practice most consecutive elements are in nearby nodes, which on average
385 reduces the number of traversals across nodes, resulting in considerably lower amortised conversion
386 cost.

387 5.3. Practical Considerations

388 Since the conversion algorithms given in Figures 4 and 5 are quite straightforward, the \mathcal{O} notation
389 does not hide any large constant factors. For COO/CSC conversions the cost is $\mathcal{O}(N)$, while for
390 CSC/RBT conversions the worst-case cost is $\mathcal{O}(N \cdot \log N)$. In contrast, many mathematical operations
391 on sparse matrices have much higher computational cost than the conversion algorithms. Even

392 simply adding two sparse matrices can be much more expensive than a conversion. Although the
 393 addition operation still takes $\mathcal{O}(N)$ time (assuming N is identical for both matrices), there is a lot of
 394 hidden constant overhead, since the sparsity pattern of the resulting matrix must be computed first [4].
 395 A similar situation applies for multiplication of two sparse matrices, which for square matrices takes
 396 $\mathcal{O}(N + n_cols)$ time [27], but in practice tends to be much slower due to the many passes and extra
 397 overhead of computing the output sparsity structure [25].

398 Sparse matrix factorisations are much more expensive, meaning that any conversion overhead is
 399 essentially negligible. A sparse LU factorisation is superlinear [28] as well as other factorisations like
 400 the Cholesky factorisation, which costs $\mathcal{O}(n_cols^{3/2})$ time [29]. Other factorisations and higher-level
 401 operations exhibit similar complexity characteristics. Given this, the cost of format conversions is
 402 heavily outweighed by the user convenience that they allow.

403 6. Empirical Evaluation

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

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

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

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

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

439 Figure 7 shows the wall-clock time taken to calculate the expressions $\text{trace}(A.t()*B)$ and
 440 $\text{diagmat}(A+B)$, with and without the aid of the automatic template-based optimisation of compound




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 .

441 expression described in Section 3. For both expressions, employing expression optimisation leads to
 442 considerable reduction in the wall-clock time. As the density increases (ie., more non-zero elements),
 443 more time is saved via expression optimisation.

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

450 7. Conclusion

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

459 The source code for the sparse matrix class and its associated functions is included in recent
460 releases of the cross-platform and open-source Armadillo linear algebra library [3], available from
461 <http://arma.sourceforge.net>. The code is provided under the permissive Apache 2.0 license [30],
462 allowing unencumbered use in both open-source and proprietary projects (eg. product development).

463 The sparse matrix class has already been successfully used in open-source projects such as
464 the *mlpack* library for machine learning [31], and the *ensmallen* library for mathematical function
465 optimisation [32]. In both cases the sparse matrix class is used to allow various algorithms to be run on
466 either sparse or dense datasets. Furthermore, bi-directional bindings for the class are provided to the R
467 environment via the Rcpp bridge [33]. Avenues for further exploration include expanding the hybrid
468 storage framework with more sparse matrix formats [4,9] in order to provide speedups for specialised
469 use cases.

470 **Author Contributions:** The authors (C.S. and R.C.) contributed equally in all aspects.

471 **Funding:** This research received no external funding.

472 **Acknowledgments:** The authors would like to thank their colleagues at the University of Queensland (Ian Hayes,
473 George Havas, Arnold Wiliem) and Data61/CSIRO (Dan Pagendam, Josh Bowden, Regis Riveret) for discussions
474 leading to the improvements of this article.

475 **Conflicts of Interest:** The authors declare no conflict of interest.

476 References

- 477 1. Nunez-Iglesias, J.; van der Walt, S.; Dashnow, H. *Elegant SciPy: The Art of Scientific Python*; O'Reilly Media,
478 2017.
- 479 2. Stroustrup, B. *The C++ Programming Language*, 4th ed.; Addison-Wesley, 2013.
- 480 3. Sanderson, C.; Curtin, R. Armadillo: a template-based C++ library for linear algebra. *Journal of Open Source
481 Software* **2016**, *1*, 26.
- 482 4. Saad, Y. SPARSKIT: A basic tool kit for sparse matrix computations. Technical Report NASA-CR-185876,
483 NASA Ames Research Center, 1990.
- 484 5. Eaton, J.W.; Bateman, D.; Hauberg, S.; Wehbring, R. *GNU Octave 4.2 Reference Manual*; Samurai Media
485 Limited, 2017.
- 486 6. Davis, T.A.; Rajamanickam, S.; Sid-Lakhdar, W.M. A survey of direct methods for sparse linear systems.
487 *Acta Numerica* **2016**, *25*, 383–566.
- 488 7. MathWorks. MATLAB Documentation - Accessing Sparse Matrices. [https://www.mathworks.com/help/
489 matlab/math/accessing-sparse-matrices.html](https://www.mathworks.com/help/matlab/math/accessing-sparse-matrices.html), 2018.
- 490 8. Iglberger, K.; Hager, G.; Treibig, J.; Rde, U. Expression templates revisited: a performance analysis of
491 current methodologies. *SIAM Journal on Scientific Computing* **2012**, *34*, C42–C69.
- 492 9. Duff, I.S.; Erisman, A.M.; Reid, J.K. *Direct methods for sparse matrices*, 2nd ed.; Oxford University Press,
493 2017.
- 494 10. Liniker, P.; Beckmann, O.; Kelly, P.H. Delayed Evaluation, Self-optimising Software Components as a
495 Programming Model. European Conference on Parallel Processing - Euro-Par 2002. Lecture Notes in
496 Computer Science (LNCS), 2002, Vol. 2400, pp. 666–673.
- 497 11. Abrahams, D.; Gurtovoy, A. *C++ Template Metaprogramming: Concepts, Tools, and Techniques from Boost and
498 Beyond*; Addison-Wesley Professional, 2004.
- 499 12. Vandevoorde, D.; Josuttis, N.M. *C++ Templates: The Complete Guide*, 2nd ed.; Addison-Wesley, 2017.
- 500 13. Sanderson, C.; Curtin, R. A User-Friendly Hybrid Sparse Matrix Class in C++. Mathematical Software -
501 ICMS 2018. Lecture Notes in Computer Science (LNCS), 2018, Vol. 10931, pp. 422–430.
- 502 14. Lehoucq, R.B.; Sorensen, D.C.; Yang, C. *ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems
503 with Implicitly Restarted Arnoldi Methods*; SIAM, 1998.
- 504 15. Li, X.S. An overview of SuperLU: Algorithms, implementation, and user interface. *ACM Transactions on
505 Mathematical Software (TOMS)* **2005**, *31*, 302–325.
- 506 16. Mernik, M.; Heering, J.; Sloane, A.M. When and how to develop domain-specific languages. *ACM
507 Computing Surveys* **2005**, *37*, 316–344.
- 508 17. Scherr, M.; Chiba, S. Almost first-class language embedding: taming staged embedded DSLs. ACM
509 SIGPLAN International Conference on Generative Programming: Concepts and Experiences, 2015, pp.
510 21–30.
- 511 18. Veldhuizen, T.L. C++ Templates as Partial Evaluation. ACM SIGPLAN Workshop on Partial Evaluation
512 and Semantics-Based Program Manipulation, 1999, pp. 13–18.
- 513 19. Vandenberghe, L.; Boyd, S. Semidefinite Programming. *SIAM Review* **1996**, *38*, 49–95.
- 514 20. Boumal, N.; Voroninski, V.; Bandeira, A. The non-convex Burer-Monteiro approach works on smooth
515 semidefinite programs. Advances in Neural Information Processing Systems, 2016, pp. 2757–2765.
- 516 21. El Ghaoui, L.; Le Bret, H. Robust solutions to least-squares problems with uncertain data. *SIAM Journal on
517 Matrix Analysis and Applications* **1997**, *18*, 1035–1064.
- 518 22. Lanckriet, G.R.; Cristianini, N.; Bartlett, P.; Ghaoui, L.E.; Jordan, M.I. Learning the Kernel Matrix with
519 Semidefinite Programming. *Journal of Machine Learning Research* **2004**, *5*, 27–72.

- 520 23. Mittal, S. A Survey of Recent Prefetching Techniques for Processor Caches. *ACM Computing Surveys* **2016**,
521 49, 35:1–35:35.
- 522 24. Cormen, T.H.; Leiserson, C.E.; Rivest, R.L.; Stein, C. *Introduction to Algorithms*, 3rd ed.; MIT Press, 2009.
- 523 25. Bank, R.E.; Douglas, C.C. Sparse matrix multiplication package (SMMP). *Advances in Computational*
524 *Mathematics* **1993**, 1, 127–137.
- 525 26. Anderson, E.; Bai, Z.; Bischof, C.; Blackford, S.; Demmel, J.; Dongarra, J.; Du Croz, J.; Greenbaum, A.;
526 Hammarling, S.; McKenney, A.; Sorensen, D. *LAPACK Users' Guide*; SIAM, 1999.
- 527 27. Davis, T.A. *Direct methods for sparse linear systems*; SIAM, 2006.
- 528 28. Gilbert, J.R.; Li, X.S.; Ng, E.G.; Peyton, B.W. Computing row and column counts for sparse QR and LU
529 factorization. *BIT Numerical Mathematics* **2001**, 41, 693–710.
- 530 29. George, A.; Ng, E. On the complexity of sparse QR and LU factorization of finite-element matrices. *SIAM*
531 *Journal on Scientific and Statistical Computing* **1988**, 9, 849–861.
- 532 30. St. Laurent, A. *Understanding Open Source and Free Software Licensing*; O'Reilly Media, 2008.
- 533 31. Curtin, R.; Edel, M.; Lozhnikov, M.; Mentekidis, Y.; Ghaisas, S.; Zhang, S. mlpack 3: a fast, flexible machine
534 learning library. *Journal of Open Source Software* **2018**, 3, 726.
- 535 32. Bhardwaj, S.; Curtin, R.; Edel, M.; Mentekidis, Y.; Sanderson, C. ensmallen: a flexible C++ library for
536 efficient function optimization. *Workshop on Systems for ML and Open Source Software at NIPS / NeurIPS*
537 **2018**.
- 538 33. Eddelbuettel, D.; Sanderson, C. RcppArmadillo: Accelerating R with high-performance C++ linear algebra.
539 *Computational Statistics & Data Analysis* **2014**, 71, 1054–1063.

540 © 2019 by the authors. Submitted to *Math. Comput. Appl.* for possible open access
541 publication under the terms and conditions of the Creative Commons Attribution (CC BY) license
542 (<http://creativecommons.org/licenses/by/4.0/>).