

# Multivariate Normal Log-likelihoods in the **mvtnorm** Package <sup>1</sup>

Torsten Hothorn

Version 1.3-0

<sup>1</sup>Please cite this document as: Torsten Hothorn (2024) Multivariate Normal Log-likelihoods in the **mvtnorm** Package. R package vignette version 1.3-0, URL [DOI:10.32614/CRAN.package.mvtnorm](https://doi.org/10.32614/CRAN.package.mvtnorm).

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Lower Triangular Matrices</b>	<b>2</b>
2.1	Multiple Lower Triangular Matrices . . . . .	3
2.2	Printing . . . . .	10
2.3	Reordering . . . . .	11
2.4	Subsetting . . . . .	12
2.5	Diagonal Elements . . . . .	18
2.6	Multiplication . . . . .	22
2.7	Solving Linear Systems . . . . .	27
2.8	Log-determinants . . . . .	32
2.9	Crossproducts . . . . .	34
2.10	Cholesky Factorisation . . . . .	38
2.11	Kronecker Products . . . . .	40
2.12	Convenience Functions . . . . .	45
2.13	Marginal and Conditional Normal Distributions . . . . .	51
2.14	Continuous Log-likelihoods . . . . .	56
2.15	Application Example . . . . .	61
<b>3</b>	<b>Multivariate Normal Log-likelihoods</b>	<b>63</b>
3.1	Algorithm . . . . .	64
3.2	Score Function . . . . .	76
<b>4</b>	<b>Maximum-likelihood Example</b>	<b>89</b>
<b>5</b>	<b>Continuous-discrete Likelihoods</b>	<b>98</b>
<b>6</b>	<b>Unstructured Gaussian Copula Estimation</b>	<b>108</b>
<b>7</b>	<b>(Experimental) User Interface</b>	<b>114</b>
<b>8</b>	<b>Package Infrastructure</b>	<b>130</b>

# Licence

Copyright (C) 2022– Torsten Hothorn

This file is part of the **mvtnorm** R add-on package.

**mvtnorm** is free software: you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation, version 2.

**mvtnorm** is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with **mvtnorm**. If not, see <<http://www.gnu.org/licenses/>>.

# Chapter 1

## Introduction

This document describes an implementation of [Genz \(1992\)](#) and, partially, of [Genz and Bretz \(2002\)](#), for the evaluation of  $N$  multivariate  $J$ -dimensional normal probabilities

$$p_i(\mathbf{C}_i \mid \mathbf{a}_i, \mathbf{b}_i) = \mathbb{P}(\mathbf{a}_i < \mathbf{Y}_i \leq \mathbf{b}_i \mid \mathbf{C}_i) = (2\pi)^{-\frac{J}{2}} \det(\mathbf{C}_i)^{-1} \int_{\mathbf{a}_i}^{\mathbf{b}_i} \exp\left(-\frac{1}{2} \mathbf{y}^\top \mathbf{C}_i^{-1} \mathbf{y}\right) d\mathbf{y} \quad (1.1)$$

where  $\mathbf{a}_i = (a_1^{(i)}, \dots, a_J^{(i)})^\top \in \mathbb{R}^J$  and  $\mathbf{b}_i = (b_1^{(i)}, \dots, b_J^{(i)})^\top \in \mathbb{R}^J$  are integration limits,  $\mathbf{C}_i = (c_{jj}^{(i)}) \in \mathbb{R}^{J \times J}$  is a lower triangular matrix with  $c_{jj}^{(i)} = 0$  for  $1 \leq j < J < J$ , and thus  $\mathbf{Y}_i \sim \mathbb{N}_J(\mathbf{0}_J, \mathbf{C}_i \mathbf{C}_i^\top)$  for  $i = 1, \dots, N$ .

One application of these integrals is the estimation of the Cholesky factor  $\mathbf{C}$  of a  $J$ -dimensional normal distribution based on  $N$  interval-censored observations  $\mathbf{Y}_1, \dots, \mathbf{Y}_N$  (encoded by  $\mathbf{a}$  and  $\mathbf{b}$ ) via maximum-likelihood

$$\hat{\mathbf{C}} = \operatorname{argmax}_{\mathbf{C}} \sum_{i=1}^N \log(p_i(\mathbf{C} \mid \mathbf{a}_i, \mathbf{b}_i)).$$

In other applications, the Cholesky factor might also depend on  $i$  in some structured way.

Function `pmvnorm` in package `mvtnorm` computes  $p_i$  based on the covariance matrix  $\mathbf{C}_i \mathbf{C}_i^\top$ . However, the Cholesky factor  $\mathbf{C}_i$  of the given covariance matrix is computed in FORTRAN first each time this function is called. Function `pmvnorm` is not vectorised over  $i = 1, \dots, N$  and thus separate calls to this function are necessary in order to compute likelihood contributions.

The implementation described here is a re-implementation (in R and C) of Alan Genz' original FORTRAN code, focusing on efficient computation of the log-likelihood  $\sum_{i=1}^N \log(p_i)$  and the corresponding score function.

The document first describes a class and some useful methods for dealing with multiple lower triangular matrices  $\mathbf{C}_i, i = 1, \dots, N$  in Chapter 2. The multivariate normal log-likelihood, and the corresponding score function, is implemented as outlined in Chapter 3. An example demonstrating maximum-likelihood estimation of Cholesky factors in the presence of interval-censored observations is discussed in Chapter 4. We use the technology developed here to implement the log-likelihood and score function for situations where some variables have been observed exactly and others only in form of interval-censoring in Chapter 5 and for nonparametric maximum-likelihood estimation in unstructured Gaussian copulae in Chapter 6. An attempt to provide useRs with a simple and (hopefully) bullet proof interface is documented in Chapter 7.

## Chapter 2

# Lower Triangular Matrices

"ltMatrices.R" 2≡

```
⟨ R Header 131 ⟩  
⟨ ltMatrices 6a ⟩  
⟨ syMatrices 6b ⟩  
⟨ dim ltMatrices 6c ⟩  
⟨ dimnames ltMatrices 7a ⟩  
⟨ names ltMatrices 7b ⟩  
⟨ is.ltMatrices 7c ⟩  
⟨ as.ltMatrices 115b ⟩  
⟨ print ltMatrices 11 ⟩  
⟨ reorder ltMatrices 12 ⟩  
⟨ subset ltMatrices 14 ⟩  
⟨ lower triangular elements 17 ⟩  
⟨ diagonals ltMatrices 19 ⟩  
⟨ diagonal matrix 22 ⟩  
⟨ mult ltMatrices 23a ⟩  
⟨ mult syMatrices 27 ⟩  
⟨ solve ltMatrices 31 ⟩  
⟨ logdet ltMatrices 33b ⟩  
⟨ tcrossprod ltMatrices 37 ⟩  
⟨ crossprod ltMatrices 38 ⟩  
⟨ chol syMatrices 39 ⟩  
⟨ add diagonal elements 20 ⟩  
⟨ assign diagonal elements 21 ⟩  
⟨ kronecker vec trick 44 ⟩  
⟨ convenience functions 48 ⟩  
⟨ aperm 51a, ... ⟩  
⟨ marginal 52b ⟩  
⟨ conditional 55 ⟩  
⟨ check obs 57b ⟩  
⟨ colSumsdnorm ltMatrices 58b ⟩  
◇
```

```

"ltMatrices.c" 3≡

< C Header 132 >
#ifndef USE_FC_LEN_T
# define USE_FC_LEN_T
#endif
#include <Rconfig.h>
#include <R_ext/Lapack.h> /* for dtptri */
#ifndef FCONE
# define FCONE
#endif
#include <R.h>
#include <Rmath.h>
#include <Rinternals.h>
#include <Rdefines.h>
< colSumsdnorm 58a >
< solve 29 >
< solve C 30 >
< logdet 33a >
< tcrossprod 36 >
< mult 24b >
< mult transpose 26 >
< chol 40 >
< vec trick 42a >
◇

```

We first define and implement infrastructure for dealing with multiple lower triangular matrices  $\mathbf{C}_i \in \mathbb{R}^{J \times J}$  for  $i = 1, \dots, N$ . We note that each such matrix  $\mathbf{C}$  can be stored in a vector of length  $J(J+1)/2$ . If all diagonal elements are one (that is,  $c_{jj}^{(i)} \equiv 1, j = 1, \dots, J$ ), the length of this vector is  $J(J-1)/2$ .

## 2.1 Multiple Lower Triangular Matrices

We can store  $N$  such matrices in an  $J(J+1)/2 \times N$  matrix (`diag = TRUE`) or, for `diag = FALSE`, in an  $J(J-1)/2 \times N$  matrix.

Each vector might define the corresponding lower triangular matrix either in row or column-major order:

$$\begin{aligned}
\mathbf{C} &= \begin{pmatrix} c_{11} & & & 0 \\ c_{21} & c_{22} & & \\ c_{31} & c_{32} & c_{33} & \\ \vdots & \vdots & & \ddots \\ c_{J1} & c_{J2} & \dots & c_{JJ} \end{pmatrix} \text{matrix indexing} \\
&= \begin{pmatrix} c_1 & & & 0 \\ c_2 & c_{J+1} & & \\ c_3 & c_{J+2} & c_{2J} & \\ \vdots & \vdots & & \ddots \\ c_J & c_{2J-1} & \dots & c_{J(J+1)/2} \end{pmatrix} \text{column-major, byrow = FALSE} \\
&= \begin{pmatrix} & c_1 & & & & 0 \\ & c_2 & & c_3 & & \\ & c_4 & & c_5 & c_6 & \\ & \vdots & & \vdots & & \ddots \\ c_{J((J+1)/2-1)+1} & c_{J((J+1)/2-1)+2} & \dots & & c_{J(J+1)/2} \end{pmatrix} \text{row-major, byrow = TRUE}
\end{aligned}$$

Based on some matrix `object`, the dimension  $J$  is computed and checked as

$\langle ltMatrices \ dim \ 4 \rangle \equiv$

```

J <- floor((1 + sqrt(1 + 4 * 2 * nrow(object))) / 2 - diag)
if (nrow(object) != J * (J - 1) / 2 + diag * J)
  stop("Dimension of object does not correspond to lower
       triangular part of a square matrix")

```

◇

Fragment referenced in [6a](#).

Typically the  $J$  dimensions are associated with names, and we therefore compute identifiers for the vector elements in either column- or row-major order on request (for later printing)

$\langle \text{ltMatrices names 5a} \rangle \equiv$

```

nonames <- FALSE
if (!isTRUE(names)) {
  if (is.character(names))
    stopifnot(is.character(names) &&
              length(unique(names)) == J)
  else
    nonames <- TRUE
} else {
  names <- as.character(1:J)
}

if (!nonames) {
  L1 <- matrix(names, nrow = J, ncol = J)
  L2 <- matrix(names, nrow = J, ncol = J, byrow = TRUE)
  L <- matrix(paste(L1, L2, sep = "."), nrow = J, ncol = J)
  if (byrow)
    rownames(object) <- t(L)[upper.tri(L, diag = diag)]
  else
    rownames(object) <- L[lower.tri(L, diag = diag)]
} # else {      ### add later
# warning("ltMatrices objects should be properly named")
# }
◇

```

Fragment referenced in [6a](#).

If `object` is already a classed object representing lower triangular matrices (we will use the class name `ltMatrices`), we might want to change the storage form from row- to column-major or the other way round.

$\langle \text{ltMatrices input 5b} \rangle \equiv$

```

if (is.ltMatrices(object)) {
  cls <- class(object)          ### keep inheriting classes
  ret <- .reorder(object, byrow = byrow)
  class(ret) <- class(object)
  return(ret)
}
◇

```

Fragment referenced in [6a](#).

The constructor essentially attaches attributes to a matrix `object`, possibly after some reordering / transposing

$\langle \text{ltMatrices 6a} \rangle \equiv$

```
ltMatrices <- function(object, diag = FALSE, byrow = FALSE, names = TRUE) {
  if (!is.matrix(object))
    object <- matrix(object, ncol = 1L)

   $\langle \text{ltMatrices input 5b} \rangle$ 
   $\langle \text{ltMatrices dim 4} \rangle$ 
   $\langle \text{ltMatrices names 5a} \rangle$ 

  attr(object, "J")      <- J
  attr(object, "diag")   <- diag
  attr(object, "byrow")  <- byrow
  attr(object, "rcnames") <- names

  class(object) <- c("ltMatrices", class(object))
  object
}
◇
```

Fragment referenced in 2.

For the sake of completeness, we also add a constructor for multiple symmetric matrices

$\langle \text{syMatrices 6b} \rangle \equiv$

```
as.syMatrices <- function(x) {
  if (is.syMatrices(x))
    return(x)
  x <- as.ltMatrices(x)      ### make sure "ltMatrices"
                             ### is first class
  class(x)[1L] <- "syMatrices"
  return(x)
}
syMatrices <- function(object, diag = FALSE, byrow = FALSE, names = TRUE)
  as.syMatrices(ltMatrices(object = object, diag = diag, byrow = byrow,
                             names = names))
◇
```

Fragment referenced in 2.

The dimensions of such an object are always  $N \times J \times J$  and are given by

$\langle \text{dim ltMatrices 6c} \rangle \equiv$

```
dim.ltMatrices <- function(x) {
  J <- attr(x, "J")
  return(c(attr(x, "dim")[2L], J, J)) ### ncol(unclass(x)) may trigger gc
}
dim.syMatrices <- dim.ltMatrices
◇
```

Fragment referenced in 2.

The corresponding dimnames can be extracted as

$\langle \text{dimnames } \text{ltMatrices } 7a \rangle \equiv$

```
dimnames.ltMatrices <- function(x)
  return(list(attr(x, "dimnames")[[2L]], attr(x, "rcnames"), attr(x, "rcnames")))
dimnames.syMatrices <- dimnames.ltMatrices
◇
```

Fragment referenced in 2.

The names identifying rows and columns in each  $C_i$  are

$\langle \text{names } \text{ltMatrices } 7b \rangle \equiv$

```
names.ltMatrices <- function(x) {
  return(attr(x, "dimnames")[[1L]])
}
names.syMatrices <- names.ltMatrices
◇
```

Fragment referenced in 2.

Finally, let's add two functions for checking the class and a function for coersing classes inheriting from `ltMatrices` to the latter, the same for `syMatrices`. Furthermode, `as.ltMatrices` coerces objects inheriting from `syMatrices` or `ltMatrices` to class `ltMatrices` (that is, `chol` or `invchol` is removed from the class list, unlike a call to the constructor `ltMatrices`). A `default` method is added in Chapter 7.

$\langle \text{is.ltMatrices } 7c \rangle \equiv$

```
is.ltMatrices <- function(x) inherits(x, "ltMatrices")
is.syMatrices <- function(x) inherits(x, "syMatrices")
as.ltMatrices <- function(x) UseMethod("as.ltMatrices")
as.ltMatrices.syMatrices <- function(x) {
  cls <- class(x)
  class(x) <- cls[which(cls == "syMatrices"):length(cls)]
  class(x)[1L] <- "ltMatrices"
  return(x)
}
as.ltMatrices.ltMatrices <- function(x) {
  cls <- class(x)
  class(x) <- cls[which(cls == "ltMatrices"):length(cls)]
  return(x)
}
◇
```

Fragment referenced in 2.

Let's set-up an example for illustration. Throughout this document, we will compare numerical results using

```
> chk <- function(...) stopifnot(isTRUE(all.equal(...)))
```

We start with a a simple example demonstrating how to set-up `ltMatrices` objects

```
> library("mvtnorm")
> set.seed(290875)
```

```

> N <- 4L
> J <- 5L
> rn <- paste0("C_", 1:N)
> nm <- LETTERS[1:J]
> Jn <- J * (J - 1) / 2
> ## data
> xn <- matrix(runif(N * Jn), ncol = N)
> colnames(xn) <- rn
> xd <- matrix(runif(N * (Jn + J)), ncol = N)
> colnames(xd) <- rn
> (lxn <- ltMatrices(xn, byrow = TRUE, names = nm))

```

```
, , C_1
```

	A	B	C	D	E
A	1.00000000	0.00000000	0.00000000	0.00000000	0
B	0.51236601	1.00000000	0.00000000	0.00000000	0
C	0.05847253	0.9095137	1.00000000	0.00000000	0
D	0.39448719	0.6612143	0.23352591	1.00000000	0
E	0.51647518	0.2979867	0.07517749	0.8182123	1

```
, , C_2
```

	A	B	C	D	E
A	1.00000000	0.00000000	0.00000000	0.00000000	0
B	0.8590665	1.00000000	0.00000000	0.00000000	0
C	0.3744315	0.1022684	1.00000000	0.00000000	0
D	0.1165248	0.7956529	0.8930589	1.00000000	0
E	0.1948049	0.4730419	0.2377852	0.214606	1

```
, , C_3
```

	A	B	C	D	E
A	1.00000000	0.00000000	0.00000000	0.00000000	0
B	0.4530153	1.00000000	0.00000000	0.00000000	0
C	0.9045608	0.9269936	1.00000000	0.00000000	0
D	0.4490011	0.1326375	0.4153967	1.00000000	0
E	0.9574833	0.4917481	0.7160702	0.2938002	1

```
, , C_4
```

	A	B	C	D	E
A	1.0000000000	0.00000000	0.0000000000	0.00000000	0
B	0.4877241328	1.00000000	0.0000000000	0.00000000	0
C	0.0593045885	0.7625270	1.0000000000	0.00000000	0
D	0.0005227393	0.1995700	0.470508903	1.00000000	0
E	0.4913541358	0.2849431	0.005961103	0.8901458	1

```
> dim(lxn)
```

```
[1] 4 5 5
```

```
> dimnames(lxn)
```

```
[[1]]
```

```
[1] "C_1" "C_2" "C_3" "C_4"
```

```

[[2]]
[1] "A" "B" "C" "D" "E"

[[3]]
[1] "A" "B" "C" "D" "E"

> lxd <- ltMatrices(xd, byrow = TRUE, diag = TRUE, names = nm)
> dim(lxd)

[1] 4 5 5

> dimnames(lxd)

[[1]]
[1] "C_1" "C_2" "C_3" "C_4"

[[2]]
[1] "A" "B" "C" "D" "E"

[[3]]
[1] "A" "B" "C" "D" "E"

> lxn <- as.syMatrices(lxn)
> lxn

, , C_1

      A          B          C          D          E
A 1.00000000 0.5123660 0.05847253 0.3944872 0.51647518
B 0.51236601 1.0000000 0.90951367 0.6612143 0.29798667
C 0.05847253 0.9095137 1.00000000 0.2335259 0.07517749
D 0.39448719 0.6612143 0.23352591 1.0000000 0.81821229
E 0.51647518 0.2979867 0.07517749 0.8182123 1.00000000

, , C_2

      A          B          C          D          E
A 1.00000000 0.8590665 0.3744315 0.1165248 0.1948049
B 0.8590665 1.0000000 0.1022684 0.7956529 0.4730419
C 0.3744315 0.1022684 1.0000000 0.8930589 0.2377852
D 0.1165248 0.7956529 0.8930589 1.0000000 0.2146060
E 0.1948049 0.4730419 0.2377852 0.2146060 1.0000000

, , C_3

      A          B          C          D          E
A 1.00000000 0.4530153 0.9045608 0.4490011 0.9574833
B 0.4530153 1.0000000 0.9269936 0.1326375 0.4917481
C 0.9045608 0.9269936 1.0000000 0.4153967 0.7160702
D 0.4490011 0.1326375 0.4153967 1.0000000 0.2938002
E 0.9574833 0.4917481 0.7160702 0.2938002 1.0000000

, , C_4

```

	A	B	C	D	E
A	1.0000000000	0.4877241	0.059304588	0.0005227393	0.491354136
B	0.4877241328	1.0000000	0.762527028	0.1995699527	0.284943077
C	0.0593045885	0.7625270	1.000000000	0.4705089033	0.005961103
D	0.0005227393	0.1995700	0.470508903	1.000000000	0.890145786
E	0.4913541358	0.2849431	0.005961103	0.8901457863	1.000000000

## 2.2 Printing

For pretty printing, we coerce objects of class `ltMatrices` to `array`. The method has a logical argument called `symmetric`, forcing the lower triangular matrix to be interpreted as a symmetric matrix.

$\langle \text{extract slots 10} \rangle \equiv$

```
diag <- attr(x, "diag")
byrow <- attr(x, "byrow")
d <- dim(x)
J <- d[2L]
dn <- dimnames(x)
◇
```

Fragment referenced in [11](#), [12](#), [13](#), [17](#), [19](#), [21](#), [23a](#), [27](#).

*< print ltMatrices 11 >*  $\equiv$

```
as.array.ltMatrices <- function(x, symmetric = FALSE, ...) {  
  < extract slots 10 >  
  x <- unclass(x)  
  
  L <- matrix(1L, nrow = J, ncol = J)  
  diag(L) <- 2L  
  if (byrow) {  
    L[upper.tri(L, diag = diag)] <- floor(2L + 1:(J * (J - 1) / 2L + diag * J))  
    L <- t(L)  
  } else {  
    L[lower.tri(L, diag = diag)] <- floor(2L + 1:(J * (J - 1) / 2L + diag * J))  
  }  
  if (symmetric) {  
    L[upper.tri(L)] <- 0L  
    dg <- diag(L)  
    L <- L + t(L)  
    diag(L) <- dg  
  }  
  ret <- rbind(0, 1, x)[c(L), , drop = FALSE]  
  class(ret) <- "array"  
  dim(ret) <- d[3:1]  
  dimnames(ret) <- dn[3:1]  
  return(ret)  
}  
  
as.array.syMatrices <- function(x, ...)  
  return(as.array.ltMatrices(x, symmetric = TRUE))  
  
print.ltMatrices <- function(x, ...)  
  print(as.array(x))  
  
print.syMatrices <- function(x, ...)  
  print(as.array(x))  
◇
```

Fragment referenced in [2](#).

Symmetric matrices are represented by lower triangular matrix objects, but we change the class from `ltMatrices` to `syMatrices` (which disables all functionality except printing and coercion to arrays).

## 2.3 Reordering

It is sometimes convenient to have access to lower triangular matrices in either column- or row-major order and this little helper function switches between the two forms

$\langle \text{reorder } \text{ltMatrices } 12 \rangle \equiv$

```
.reorder <- function(x, byrow = FALSE) {

  stopifnot(is.ltMatrices(x))
  if (attr(x, "byrow") == byrow) return(x)

   $\langle \text{extract slots } 10 \rangle$ 

  x <- unclass(x)

  rL <- cL <- diag(0, nrow = J)
  rL[lower.tri(rL, diag = diag)] <- cL[upper.tri(cL, diag = diag)] <- 1:nrow(x)
  cL <- t(cL)
  if (byrow) ### row -> col order
    return(ltMatrices(x[cL[lower.tri(cL, diag = diag)], , drop = FALSE],
                      diag = diag, byrow = FALSE, names = dn[[2L]]))
  ### col -> row order
  return(ltMatrices(x[t(rL)[upper.tri(rL, diag = diag)], , drop = FALSE],
                    diag = diag, byrow = TRUE, names = dn[[2L]]))
}
 $\diamond$ 
```

Fragment referenced in 2.

We can check if this works by switching back and forth between column-major and row-major order

```
> ## constructor + .reorder + as.array
> a <- as.array(ltMatrices(xn, byrow = TRUE))
> b <- as.array(ltMatrices(ltMatrices(xn, byrow = TRUE),
+                          byrow = FALSE))
> chk(a, b)
> a <- as.array(ltMatrices(xn, byrow = FALSE))
> b <- as.array(ltMatrices(ltMatrices(xn, byrow = FALSE),
+                          byrow = TRUE))
> chk(a, b)
> a <- as.array(ltMatrices(xd, byrow = TRUE, diag = TRUE))
> b <- as.array(ltMatrices(ltMatrices(xd, byrow = TRUE, diag = TRUE),
+                          byrow = FALSE))
> chk(a, b)
> a <- as.array(ltMatrices(xd, byrow = FALSE, diag = TRUE))
> b <- as.array(ltMatrices(ltMatrices(xd, byrow = FALSE, diag = TRUE),
+                          byrow = TRUE))
> chk(a, b)
```

## 2.4 Subsetting

We might want to select subsets of observations  $i \in \{1, \dots, N\}$  or rows/columns  $j \in \{1, \dots, J\}$  of the corresponding matrices  $\mathbf{C}_i$ .

`<.subset ltMatrices 13> ≡`

```
.subset_ltMatrices <- function(x, i, j, ..., drop = FALSE) {

  if (drop) warning("argument drop is ignored")
  if (missing(i) && missing(j)) return(x)

  < extract slots 10 >

  x <- unclass(x)

  if (!missing(j)) {

    if (is.character(j)) {
      stopifnot(all(j %in% dn[[2L]]))
      j <- match(j, dn[[2L]])
    }

    j <- (1:J)[j] ### get rid of negative indices

    if (length(j) == 1L && !diag) {
      return(ltMatrices(matrix(1, ncol = ncol(x), nrow = 1), diag = TRUE,
                               byrow = byrow, names = dn[[2L]][j]))
    }
    L <- diag(0L, nrow = J)
    Jp <- sum(upper.tri(L, diag = diag))
    if (byrow) {
      L[upper.tri(L, diag = diag)] <- 1:Jp
      L <- L + t(L)
      diag(L) <- diag(L) / 2
      L <- L[j, j, drop = FALSE]
      L <- L[upper.tri(L, diag = diag)]
    } else {
      L[lower.tri(L, diag = diag)] <- 1:Jp
      L <- L + t(L)
      diag(L) <- diag(L) / 2
      L <- L[j, j, drop = FALSE]
      L <- L[lower.tri(L, diag = diag)]
    }
    if (missing(i)) {
      return(ltMatrices(x[c(L), , drop = FALSE], diag = diag,
                        byrow = byrow, names = dn[[2L]][j]))
    }
    return(ltMatrices(x[c(L), i, drop = FALSE], diag = diag,
                      byrow = byrow, names = dn[[2L]][j]))
  }
  return(ltMatrices(x[, i, drop = FALSE], diag = diag,
                    byrow = byrow, names = dn[[2L]]))
}
◇
```

Fragment referenced in [14](#).

$\langle \text{subset ltMatrices 14} \rangle \equiv$

```

<.subset ltMatrices 13>
### if j is not ordered, result is not a lower triangular matrix
".ltMatrices" <- function(x, i, j, ..., drop = FALSE) {
  if (!missing(j)) {
    if (is.character(j)) {
      stopifnot(all(j %in% dimnames(x)[[2L]]))
      j <- match(j, dimnames(x)[[2L]])
    }
    if (all(j > 0)) {
      if (any(diff(j) < 0)) stop("invalid subset argument j")
    }
  }

  return(.subset_ltMatrices(x = x, i = i, j = j, ..., drop = drop))
}

".syMatrices" <- function(x, i, j, ..., drop = FALSE) {
  x <- as.syMatrices(x)
  ret <- .subset_ltMatrices(x = x, i = i, j = j, ..., drop = drop)
  class(ret)[1L] <- "syMatrices"
  ret
}

```

Fragment referenced in [2](#).

We check if this works by first subsetting the `ltMatrices` object. Second, we coerce the object to an array and do the subset for the latter object. Both results must agree.

```

> ## subset
> a <- as.array(ltMatrices(xn, byrow = FALSE, names = nm)[i, j])
> b <- as.array(ltMatrices(xn, byrow = FALSE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xn, byrow = TRUE, names = nm)[i, j])
> b <- as.array(ltMatrices(xn, byrow = TRUE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xd, byrow = FALSE,
+                       diag = TRUE, names = nm)[i, j])
> b <- as.array(ltMatrices(xd, byrow = FALSE,
+                       diag = TRUE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xd, byrow = TRUE, diag = TRUE,
+                       names = nm)[i, j])
> b <- as.array(ltMatrices(xd, byrow = TRUE, diag = TRUE,
+                       names = nm))[j, j, i]
> chk(a, b)

```

We start with both indices being positive integers

```

> i <- colnames(xn)[1:2]
> j <- 2:4
> ## subset
> a <- as.array(ltMatrices(xn, byrow = FALSE, names = nm)[i, j])
> b <- as.array(ltMatrices(xn, byrow = FALSE, names = nm))[j, j, i]

```

```

> chk(a, b)
> a <- as.array(ltMatrices(xn, byrow = TRUE, names = nm)[i, j])
> b <- as.array(ltMatrices(xn, byrow = TRUE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xd, byrow = FALSE,
+                       diag = TRUE, names = nm)[i, j])
> b <- as.array(ltMatrices(xd, byrow = FALSE,
+                       diag = TRUE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xd, byrow = TRUE, diag = TRUE,
+                       names = nm)[i, j])
> b <- as.array(ltMatrices(xd, byrow = TRUE, diag = TRUE,
+                       names = nm))[j, j, i]
> chk(a, b)

```

proceed with characters

```

> i <- 1:2
> j <- nm[2:4]
> ## subset
> a <- as.array(ltMatrices(xn, byrow = FALSE, names = nm)[i, j])
> b <- as.array(ltMatrices(xn, byrow = FALSE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xn, byrow = TRUE, names = nm)[i, j])
> b <- as.array(ltMatrices(xn, byrow = TRUE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xd, byrow = FALSE,
+                       diag = TRUE, names = nm)[i, j])
> b <- as.array(ltMatrices(xd, byrow = FALSE,
+                       diag = TRUE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xd, byrow = TRUE, diag = TRUE,
+                       names = nm)[i, j])
> b <- as.array(ltMatrices(xd, byrow = TRUE, diag = TRUE,
+                       names = nm))[j, j, i]
> chk(a, b)

```

a different subset

```

> j <- c(1, 3, 5)
> ## subset
> a <- as.array(ltMatrices(xn, byrow = FALSE, names = nm)[i, j])
> b <- as.array(ltMatrices(xn, byrow = FALSE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xn, byrow = TRUE, names = nm)[i, j])
> b <- as.array(ltMatrices(xn, byrow = TRUE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xd, byrow = FALSE,
+                       diag = TRUE, names = nm)[i, j])
> b <- as.array(ltMatrices(xd, byrow = FALSE,
+                       diag = TRUE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xd, byrow = TRUE, diag = TRUE,
+                       names = nm)[i, j])
> b <- as.array(ltMatrices(xd, byrow = TRUE, diag = TRUE,

```

```

+                               names = nm))[j, j, i]
> chk(a, b)

    and characters again

> j <- nm[c(1, 3, 5)]
> ## subset
> a <- as.array(ltMatrices(xn, byrow = FALSE, names = nm)[i, j])
> b <- as.array(ltMatrices(xn, byrow = FALSE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xn, byrow = TRUE, names = nm)[i, j])
> b <- as.array(ltMatrices(xn, byrow = TRUE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xd, byrow = FALSE,
+                               diag = TRUE, names = nm)[i, j])
> b <- as.array(ltMatrices(xd, byrow = FALSE,
+                               diag = TRUE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xd, byrow = TRUE, diag = TRUE,
+                               names = nm)[i, j])
> b <- as.array(ltMatrices(xd, byrow = TRUE, diag = TRUE,
+                               names = nm))[j, j, i]
> chk(a, b)

```

and finally with with negative subsets

```

> j <- -c(1, 3, 5)
> ## subset
> a <- as.array(ltMatrices(xn, byrow = FALSE, names = nm)[i, j])
> b <- as.array(ltMatrices(xn, byrow = FALSE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xn, byrow = TRUE, names = nm)[i, j])
> b <- as.array(ltMatrices(xn, byrow = TRUE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xd, byrow = FALSE,
+                               diag = TRUE, names = nm)[i, j])
> b <- as.array(ltMatrices(xd, byrow = FALSE,
+                               diag = TRUE, names = nm))[j, j, i]
> chk(a, b)
> a <- as.array(ltMatrices(xd, byrow = TRUE, diag = TRUE,
+                               names = nm)[i, j])
> b <- as.array(ltMatrices(xd, byrow = TRUE, diag = TRUE,
+                               names = nm))[j, j, i]
> chk(a, b)

```

and with non-increasing argument j (this won't work for lower triangular matrices, only for symmetric matrices)

```

> ## subset
> j <- nm[sample(1:J)]
> ltM <- ltMatrices(xn, byrow = FALSE, names = nm)
> try(ltM[i, j])
> ltM <- as.syMatrices(ltM)
> a <- as.array(ltM[i, j])
> b <- as.array(ltM)[j, j, i]
> chk(a, b)

```

Extracting the lower triangular elements from an `ltMatrices` object (or from an object of class `syMatrices`) returns a matrix with  $N$  columns, undoing the effect of `ltMatrices`. Note that ordering of the rows of this matrix depend on the `byrow` attribute of `x`, unless the `byrow` to this function is used to overwrite it explicitly

*(lower triangular elements 17)*  $\equiv$

```
Lower_tri <- function(x, diag = FALSE, byrow = attr(x, "byrow")) {

  if (is.syMatrices(x))
    x <- as.ltMatrices(x)
  adiaq <- diag
  x <- ltMatrices(x, byrow = byrow)

  (extract slots 10)

  if (diag == adiaq)
    return(unclass(x)[,,drop = FALSE]) ### remove attributes

  if (!diag && adiaq) {
    diagonals(x) <- 1
    return(unclass(x)[,,drop = FALSE]) ### remove attributes
  }

  x <- unclass(x)
  if (J == 1) {
    idx <- 1L
  } else {
    if (byrow)
      idx <- cumsum(c(1, 2:J))
    else
      idx <- cumsum(c(1, J:2))
  }
  return(x[-idx,,drop = FALSE])
}
◇
```

Fragment referenced in 2.

```
> ## J <- 4
> M <- ltMatrices(matrix(1:10, nrow = 10, ncol = 2), diag = TRUE)
> Lower_tri(M, diag = FALSE)
```

```
      [,1] [,2]
2.1      2      2
3.1      3      3
4.1      4      4
3.2      6      6
4.2      7      7
4.3      9      9
```

```
> Lower_tri(M, diag = TRUE)
```

```
      [,1] [,2]
1.1      1      1
2.1      2      2
3.1      3      3
```

```

4.1    4    4
2.2    5    5
3.2    6    6
4.2    7    7
3.3    8    8
4.3    9    9
4.4   10   10

> M <- ltMatrices(matrix(1:6, nrow = 6, ncol = 2), diag = FALSE)
> Lower_tri(M, diag = FALSE)

      [,1] [,2]
2.1      1      1
3.1      2      2
4.1      3      3
3.2      4      4
4.2      5      5
4.3      6      6

> Lower_tri(M, diag = TRUE)

      [,1] [,2]
1.1      1      1
2.1      1      1
3.1      2      2
4.1      3      3
2.2      1      1
3.2      4      4
4.2      5      5
3.3      1      1
4.3      6      6
4.4      1      1

> ## multiple symmetric matrices
> Lower_tri(invchol2cor(M))

      [,1]      [,2]
2.1 -0.7071068 -0.7071068
3.1  0.4364358  0.4364358
4.1 -0.4481107 -0.4481107
3.2 -0.9258201 -0.9258201
4.2  0.9189002  0.9189002
4.3 -0.9974149 -0.9974149

```

## 2.5 Diagonal Elements

The diagonal elements of each matrix  $\mathbf{C}_i$  can be extracted and are always returned as an  $J \times N$  matrix.

$\langle \text{diagonals ltMatrices 19} \rangle \equiv$

```
diagonals <- function(x, ...)
  UseMethod("diagonals")

diagonals.ltMatrices <- function(x, ...) {
   $\langle \text{extract slots 10} \rangle$ 

  x <- unclass(x)

  if (!diag) {
    ret <- matrix(1, nrow = J, ncol = ncol(x))
    colnames(ret) <- dn[[1L]]
    rownames(ret) <- dn[[2L]]
    return(ret)
  } else {
    if (J == 1L) return(x)
    if (byrow)
      idx <- cumsum(c(1, 2:J))
    else
      idx <- cumsum(c(1, J:2))
    ret <- x[idx, , drop = FALSE]
    rownames(ret) <- dn[[2L]]
    return(ret)
  }
}

diagonals.syMatrices <- diagonals.ltMatrices

diagonals.matrix <- function(x, ...) diag(x)
◇
```

Fragment referenced in [2](#).

```
> all(diagonals(ltMatrices(xn, byrow = TRUE)) == 1L)
[1] TRUE
```

Sometimes we need to add diagonal elements to an `ltMatrices` object which was set-up with constant  $c_{jj} = 1$  diagonal elements.

$\langle \text{add diagonal elements } 20 \rangle \equiv$

```
.adddiag <- function(x) {  
  stopifnot(is.ltMatrices(x))  
  
  if (attr(x, "diag")) return(x)  
  
  byrow_orig <- attr(x, "byrow")  
  
  x <- ltMatrices(x, byrow = FALSE)  
  
  N <- dim(x)[1L]  
  J <- dim(x)[2L]  
  nm <- dimnames(x)[[2L]]  
  
  L <- diag(J)  
  L[lower.tri(L, diag = TRUE)] <- 1:(J * (J + 1) / 2)  
  
  D <- diag(J)  
  ret <- matrix(D[lower.tri(D, diag = TRUE)],  
               nrow = J * (J + 1) / 2, ncol = N)  
  colnames(ret) <- dimnames(x)[[1L]]  
  ret[L[lower.tri(L, diag = FALSE)],] <- unclass(x)  
  
  ret <- ltMatrices(ret, diag = TRUE, byrow = FALSE, names = nm)  
  ret <- ltMatrices(ret, byrow = byrow_orig)  
  
  ret  
}
```

◇

Fragment referenced in [2](#).

$\langle \text{assign diagonal elements 21} \rangle \equiv$

```
"diagonals<-" <- function(x, value)
  UseMethod("diagonals<-")

"diagonals<-.ltMatrices" <- function(x, value) {
   $\langle \text{extract slots 10} \rangle$ 

  if (byrow)
    idx <- cumsum(c(1, 2:J))
  else
    idx <- cumsum(c(1, J:2))

  ### diagonals(x) <- NULL returns ltMatrices(..., diag = FALSE)
  if (is.null(value)) {
    if (!attr(x, "diag")) return(x)
    if (J == 1L) {
      x[] <- 1
      return(x)
    }
    return(ltMatrices(unclass(x)[-idx,,drop = FALSE], diag = FALSE,
                      byrow = byrow, names = dn[[2L]]))
  }

  x <- .adddiag(x)

  if (!is.matrix(value))
    value <- matrix(value, nrow = J, ncol = d[1L])

  stopifnot(is.matrix(value) && nrow(value) == J
            && ncol(value) == d[1L])

  if (J == 1L) {
    x[] <- value
    return(x)
  }

  x[idx, ] <- value

  return(x)
}

"diagonals<-.syMatrices" <- function(x, value) {

  x <- as.ltMatrices(x)
  diagonals(x) <- value
  class(x)[1L] <- "syMatrices"

  return(x)
}
 $\diamond$ 
```

Fragment referenced in [2](#).

```
> lxd2 <- lxn
> diagonals(lxd2) <- 1
> chk(as.array(lxd2), as.array(lxn))
```

A unit diagonal matrix is not treated as a special case but as an `ltMatrices` object with all lower triangular elements being zero

$\langle \text{diagonal matrix } 22 \rangle \equiv$

```
diagonals.integer <- function(x, ...)
  ltMatrices(rep(0, x * (x - 1) / 2), diag = FALSE, ...)
  ◇
```

Fragment referenced in [2](#).

```
> (I5 <- diagonals(5L))
```

```
, , 1
```

```
  1 2 3 4 5
1 1 0 0 0 0
2 0 1 0 0 0
3 0 0 1 0 0
4 0 0 0 1 0
5 0 0 0 0 1
```

```
> diagonals(I5) <- 1:5
```

```
> I5
```

```
, , 1
```

```
  1 2 3 4 5
1 1 0 0 0 0
2 0 2 0 0 0
3 0 0 3 0 0
4 0 0 0 4 0
5 0 0 0 0 5
```

## 2.6 Multiplication

Products  $\mathbf{C}_i \mathbf{y}_i$  or  $\mathbf{C}_i^\top \mathbf{y}_i$  with  $\mathbf{y}_i \in \mathbb{R}^J$  for  $i = 1, \dots, N$  can be computed with  $\mathbf{y}$  being an  $J \times N$  matrix of columns-wise stacked vectors  $(\mathbf{y}_1 \mid \mathbf{y}_2 \mid \dots \mid \mathbf{y}_N)$ . If  $\mathbf{y}$  is a single vector, it is recycled  $N$  times.

If the number of columns of a matrix  $\mathbf{y}$  is neither one nor  $N$ , we compute  $\mathbf{C}_i \mathbf{y}_j$  for all  $i = 1, \dots, N$  and  $j$ . This is dangerous but needed in [Section 2.13](#) for defining `cond_mvnorm` later on.

For  $\mathbf{C}_i \mathbf{y}_i$ , we call C code computing the product efficiently without copying data by leveraging the lower triangular structure of  $\mathbf{x} = \mathbf{C}_i$

$\langle \text{mult ltMatrices 23a} \rangle \equiv$

```
### C %*% y
Mult <- function(x, y, ...)
  UseMethod("Mult")
Mult.default <- function(x, y, transpose = FALSE, ...) {
  if (!transpose) return(x %*% y)
  return(crossprod(x, y))
}
Mult.ltMatrices <- function(x, y, transpose = FALSE, ...) {

   $\langle \text{extract slots 10} \rangle$ 

  stopifnot(is.numeric(y))
  if (!is.matrix(y)) y <- matrix(y, nrow = d[2L], ncol = d[1L])
  N <- ifelse(d[1L] == 1, ncol(y), d[1L])
  stopifnot(nrow(y) == d[2L])
  if (ncol(y) != N)
    return(sapply(1:ncol(y), function(i) Mult(x, y[,i], transpose = transpose)))

   $\langle \text{mult ltMatrices transpose 25} \rangle$ 

  x <- ltMatrices(x, byrow = TRUE)
  if (!is.double(x)) storage.mode(x) <- "double"
  if (!is.double(y)) storage.mode(y) <- "double"

  ret <- .Call(mvtnorm_R_ltMatrices_Mult, x, y, as.integer(N),
    as.integer(d[2L]), as.logical(diag))

  rownames(ret) <- dn[[2L]]
  if (length(dn[[1L]]) == N)
    colnames(ret) <- dn[[1L]]
  return(ret)
}
◇
```

Fragment referenced in [2](#).

The underlying C code assumes  $\mathbf{C}_i$  (here called C) to be in row-major order.

$\langle \text{RC input 23b} \rangle \equiv$

```
/* pointer to C matrices */
double *dC = REAL(C);
/* number of matrices */
int iN = INTEGER(N)[0];
/* dimension of matrices */
int iJ = INTEGER(J)[0];
/* C contains diagonal elements */
Rboolean Rdiag = asLogical(diag);
/* p = J * (J - 1) / 2 + diag * J */
int len = iJ * (iJ - 1) / 2 + Rdiag * iJ;
◇
```

Fragment referenced in [24b](#), [26](#), [29](#), [30](#), [33a](#), [36](#), [42a](#).

We also allow  $\mathbf{C}_i$  to be constant ( $N$  is then determined from  $\text{ncol}(y)$ ). The following fragment ensures that we only loop over  $\mathbf{C}_i$  if  $\text{dim}(x)[1L] > 1$

$\langle C \text{ length } 24a \rangle \equiv$

```
int p;
if (LENGTH(C) == len)
  /* C is constant for i = 1, ..., N */
  p = 0;
else
  /* C contains C_1, ..., C_N */
  p = len;
◇
```

Fragment referenced in [24b](#), [26](#), [29](#), [33a](#), [42a](#).

The C workhorse is now

$\langle \text{mult } 24b \rangle \equiv$

```
SEXP R_ltMatrices_Mult (SEXP C, SEXP y, SEXP N, SEXP J, SEXP diag) {

  SEXP ans;
  double *dans, *dy = REAL(y);
  int i, j, k, start;

   $\langle RC \text{ input } 23b \rangle$ 
   $\langle C \text{ length } 24a \rangle$ 

  PROTECT(ans = allocMatrix(REALSXP, iJ, iN));
  dans = REAL(ans);

  for (i = 0; i < iN; i++) {
    start = 0;
    for (j = 0; j < iJ; j++) {
      dans[j] = 0.0;
      for (k = 0; k < j; k++)
        dans[j] += dC[start + k] * dy[k];
      if (Rdiag) {
        dans[j] += dC[start + j] * dy[j];
        start += j + 1;
      } else {
        dans[j] += dy[j];
        start += j;
      }
    }
    dC += p;
    dy += iJ;
    dans += iJ;
  }
  UNPROTECT(1);
  return(ans);
}
◇
```

Fragment referenced in [3](#).

Some checks for  $C_i y_i$

```
> lxn <- ltMatrices(xn, byrow = TRUE)
> lxd <- ltMatrices(xd, byrow = TRUE, diag = TRUE)
```

```

> y <- matrix(runif(N * J), nrow = J)
> a <- Mult(lxn, y)
> A <- as.array(lxn)
> b <- do.call("rbind", lapply(1:ncol(y),
+   function(i) t(A[,i] %*% y[,i,drop = FALSE]))))
> chk(a, t(b), check.attributes = FALSE)
> a <- Mult(lxd, y)
> A <- as.array(lxd)
> b <- do.call("rbind", lapply(1:ncol(y),
+   function(i) t(A[,i] %*% y[,i,drop = FALSE]))))
> chk(a, t(b), check.attributes = FALSE)
> ### recycle C
> chk(Mult(lxn[rep(1, N),], y), Mult(lxn[1,], y), check.attributes = FALSE)
> ### recycle y
> chk(Mult(lxn, y[,1]), Mult(lxn, y[,rep(1, N)]))
> ### tcrossprod as multiplication
> i <- sample(1:N)[1]
> M <- t(as.array(lxn)[,i])
> a <- sapply(1:J, function(j) Mult(lxn[i,], M[,j,drop = FALSE]))
> rownames(a) <- colnames(a) <- dimnames(lxn)[[2L]]
> b <- as.array(Tcrossprod(lxn[i,]))[,1]
> chk(a, b, check.attributes = FALSE)

```

For  $\mathbf{C}_i^\top \mathbf{y}_i$  (transpose = TRUE), we add a dedicated C function paying attention to the lower triangular structure of  $\mathbf{x} = \mathbf{C}_i$ . This function assumes  $\mathbf{x}$  in column-major order, so we coerce this object when necessary:

$\langle \text{mult ltMatrices transpose 25} \rangle \equiv$

```

if (transpose) {
  x <- ltMatrices(x, byrow = FALSE)
  if (!is.double(x)) storage.mode(x) <- "double"
  if (!is.double(y)) storage.mode(y) <- "double"

  ret <- .Call(mvtnorm_R_ltMatrices_Mult_transpose, x, y, as.integer(N),
    as.integer(d[2L]), as.logical(diag))

  rownames(ret) <- dn[[2L]]
  if (length(dn[[1L]]) == N)
    colnames(ret) <- dn[[1L]]
  return(ret)
}

```

Fragment referenced in [23a](#).

before moving to C for the low-level computations:

$\langle \text{mult transpose 26} \rangle \equiv$

```
SEXP R_ltMatrices_Mult_transpose (SEXP C, SEXP y, SEXP N, SEXP J, SEXP diag) {

    SEXP ans;
    double *dans, *dy = REAL(y);
    int i, j, k, start;

     $\langle RC \text{ input 23b} \rangle$ 
     $\langle C \text{ length 24a} \rangle$ 

    PROTECT(ans = allocMatrix(REALSXP, iJ, iN));
    dans = REAL(ans);

    for (i = 0; i < iN; i++) {
        start = 0;
        for (j = 0; j < iJ; j++) {
            dans[j] = 0.0;
            if (Rdiag) {
                dans[j] += dC[start] * dy[j];
                start++;
            } else {
                dans[j] += dy[j];
            }
            for (k = 0; k < (iJ - j - 1); k++)
                dans[j] += dC[start + k] * dy[j + k + 1];
            start += iJ - j - 1;
        }
        dC += p;
        dy += iJ;
        dans += iJ;
    }
    UNPROTECT(1);
    return(ans);
}
◇
```

Fragment referenced in 3.

and wrap-up with some tests for computing  $\mathbf{C}_i^\top \mathbf{y}_i$

```
> a <- Mult(lxn, y, transpose = TRUE)
> A <- as.array(lxn)
> b <- do.call("rbind", lapply(1:ncol(y),
+   function(i) t(t(A[,i]) %*% y[,i,drop = FALSE]))))
> chk(a, t(b), check.attributes = FALSE)
> a <- Mult(lxd, y, transpose = TRUE)
> A <- as.array(lxd)
> b <- do.call("rbind", lapply(1:ncol(y),
+   function(i) t(t(A[,i]) %*% y[,i,drop = FALSE]))))
> chk(a, t(b), check.attributes = FALSE)
> ### recycle C
> chk(Mult(lxn[rep(1, N)], y, transpose = TRUE),
+   Mult(lxn[1,], y, transpose = TRUE), check.attributes = FALSE)
> ### recycle y
> chk(Mult(lxn, y[,1], transpose = TRUE),
+   Mult(lxn, y[,rep(1, N)], transpose = TRUE))
```

Now we can add a `Mult` method for multiple symmetric matrices, noting that for a symmetric matrix  $\mathbf{A} = \mathbf{C} + \mathbf{C}^\top - \text{diag}(\mathbf{C})$  with lower triangular part  $\mathbf{C}$  (including the diagonal) we can compute  $\mathbf{A}\mathbf{y} = \mathbf{C}\mathbf{y} + \mathbf{C}^\top\mathbf{y} - \text{diag}(\mathbf{C})\mathbf{y}$  using `Mult` applied to the lower triangular part:

`<mult syMatrices 27> ≡`

```
Mult.syMatrices <- function(x, y, ...) {
  <extract slots 10>

  x <- as.ltMatrices(x)
  stopifnot(is.numeric(y))
  if (!is.matrix(y)) y <- matrix(y, nrow = d[2L], ncol = d[1L])
  N <- ifelse(d[1L] == 1, ncol(y), d[1L])
  stopifnot(nrow(y) == d[2L])
  stopifnot(ncol(y) == N)

  ret <- Mult(x, y) + Mult(x, y, transpose = TRUE) - y * c(diagonals(x))
  return(ret)
}
◇
```

Fragment referenced in [2](#).

```
> J <- 5
> N1 <- 10
> ex <- expression({
+   C <- syMatrices(matrix(runif(N2 * J * (J + c(-1, 1)[DIAG + 1L]) / 2),
+                         ncol = N2),
+                   diag = DIAG)
+   x <- matrix(runif(N1 * J), nrow = J)
+   Ca <- as.array(C)
+   p1 <- do.call("cbind", lapply(1:N1, function(i)
+     Ca[,c(1,i)][(N2 > 1) + 1] %*% x[,i]))
+   p2 <- Mult(C, x)
+   chk(p1, p2)
+ })
> N2 <- N1
> DIAG <- TRUE
> eval(ex)
> N2 <- 1
> DIAG <- TRUE
> eval(ex)
> N2 <- 1
> DIAG <- FALSE
> eval(ex)
> N2 <- N1
> DIAG <- FALSE
> eval(ex)
```

## 2.7 Solving Linear Systems

Computing  $\mathbf{C}_i^{-1}$  or solving  $\mathbf{C}_i\mathbf{x}_i = \mathbf{y}_i$  for  $\mathbf{x}_i$  for all  $i = 1, \dots, N$  is another important task. We sometimes also need  $\mathbf{C}_i^\top\mathbf{x}_i = \mathbf{y}_i$  triggered by `transpose = TRUE`.

$\mathbf{C}$  is  $\mathbf{C}_i, i = 1, \dots, N$  in column-major order (matrix of dimension  $J(J-1)/2 + J\text{diag} \times N$ ), and  $\mathbf{y}$  is the  $J \times N$  matrix  $(\mathbf{y}_1 \mid \mathbf{y}_2 \mid \dots \mid \mathbf{y}_N)$ . This function returns the  $J \times N$  matrix  $(\mathbf{x}_1 \mid \mathbf{x}_2 \mid \dots \mid \mathbf{x}_N)$  of solutions.

If  $\mathbf{y}$  is not given,  $\mathbf{C}_i^{-1}$  is returned in the same order as the original matrix  $\mathbf{C}_i$ . If all  $\mathbf{C}_i$  have unit diagonals, so will  $\mathbf{C}_i^{-1}$ .

We start with some options for the LAPACK workhorses

$\langle \text{lapack options 28} \rangle \equiv$

```
char di, lo = 'L';
if (Rdiag) {
    /* non-unit diagonal elements */
    di = 'N';
} else {
    /* unit diagonal elements; NOTE: these diagonals is ARE always present but
       ignored in the computations */
    di = 'U';
}
◇
```

Fragment referenced in [29](#), [30](#).

and set-up a dedicated C function for computing  $\mathbf{C}_i \mathbf{x}_i = \mathbf{y}_i$

$\langle \text{solve 29} \rangle \equiv$

```
SEXP R_ltMatrices_solve (SEXP C, SEXP y, SEXP N, SEXP J, SEXP diag, SEXP transpose)
{
    SEXP ans;
    double *dans, *dy;
    int i, ONE = 1;

     $\langle RC \text{ input 23b} \rangle$ 
    /* diagonal elements are always present */
    if (!Rdiag) len += iJ;
     $\langle C \text{ length 24a} \rangle$ 
     $\langle \text{lapack options 28} \rangle$ 

    char tr = 'N';
    /* t(C) instead of C */
    Rboolean Rtranspose = asLogical(transpose);
    if (Rtranspose) {
        /* t(C) */
        tr = 'T';
    } else {
        /* C */
        tr = 'N';
    }

    dy = REAL(y);
    PROTECT(ans = allocMatrix(REALSXP, iJ, iN));
    dans = REAL(ans);
    memcpy(dans, dy, iJ * iN * sizeof(double));

    /* loop over matrices, ie columns of C / y */
    for (i = 0; i < iN; i++) {

        /* solve linear system */
        F77_CALL(dtpsv)(&lo, &tr, &di, &iJ, dC, dans, &ONE FCONE FCONE FCONE);
        dans += iJ;
        dC += p;
    }

    UNPROTECT(1);
    return(ans);
}
◇
```

Fragment referenced in 3.

and then for computing  $\mathbf{C}_i^{-1}$  explicitly

$\langle \text{solve } C \text{ 30} \rangle \equiv$

```

SEXP R_ltMatrices_solve_C (SEXP C, SEXP N, SEXP J, SEXP diag, SEXP transpose)
{
    SEXP ans;
    double *dans;
    int i, info;

     $\langle RC \text{ input 23b} \rangle$ 
    /* diagonal elements are always present */
    if (!Rdiag) len += iJ;
     $\langle \text{lapack options 28} \rangle$ 

    PROTECT(ans = allocMatrix(REALSXP, len, iN));
    dans = REAL(ans);
    memcpy(dans, dC, iN * len * sizeof(double));

    /* loop over matrices, ie columns of C / y */
    for (i = 0; i < iN; i++) {

        /* compute inverse */
        F77_CALL(dtptri)(&lo, &di, &iJ, dans, &info FCONE FCONE);
        if (info != 0)
            error("Cannot solve ltmatrices");

        dans += len;
    }

    UNPROTECT(1);
    /* note: ans always includes diagonal elements */
    return(ans);
}

```

Fragment referenced in [3](#).

with R interface

$\langle \text{solve } \text{ltMatrices } 31 \rangle \equiv$

```

solve.ltMatrices <- function(a, b, transpose = FALSE, ...) {

  byrow_orig <- attr(a, "byrow")

  x <- ltMatrices(a, byrow = FALSE)
  diag <- attr(x, "diag")
  ### dtptri and dtpsv require diagonal elements being present
  if (!diag) diagonals(x) <- diagonals(x)
  d <- dim(x)
  J <- d[2L]
  dn <- dimnames(x)
  if (!is.double(x)) storage.mode(x) <- "double"

  if (!missing(b)) {
    if (!is.matrix(b)) b <- matrix(b, nrow = J, ncol = d[1L])
    stopifnot(nrow(b) == J)
    N <- ifelse(d[1L] == 1, ncol(b), d[1L])
    stopifnot(ncol(b) == N)
    if (!is.double(b)) storage.mode(b) <- "double"
    ret <- .Call(mvtnorm_R_ltMatrices_solve, x, b,
                as.integer(N), as.integer(J), as.logical(diag),
                as.logical(transpose))
    if (d[1L] == N) {
      colnames(ret) <- dn[[1L]]
    } else {
      colnames(ret) <- colnames(b)
    }
    rownames(ret) <- dn[[2L]]
    return(ret)
  }

  if (transpose) stop("cannot compute inverse of t(a)")
  ret <- .Call(mvtnorm_R_ltMatrices_solve_C, x,
              as.integer(d[1L]), as.integer(J), as.logical(diag),
              as.logical(FALSE))
  colnames(ret) <- dn[[1L]]

  if (!diag)
    ### ret always includes diagonal elements, remove here
    ret <- ret[- cumsum(c(1, J:2)), , drop = FALSE]

  ret <- ltMatrices(ret, diag = diag, byrow = FALSE, names = dn[[2L]])
  ret <- ltMatrices(ret, byrow = byrow_orig)
  return(ret)
}

```

Fragment referenced in [2](#).

and some checks

```

> ## solve
> A <- as.array(1xn)
> a <- solve(1xn)
> a <- as.array(a)
> b <- array(apply(A, 3L, function(x) solve(x), simplify = TRUE),

```

```

+           dim = rev(dim(lxn)))
> chk(a, b, check.attributes = FALSE)
> A <- as.array(lxd)
> a <- as.array(solve(lxd))
> b <- array(apply(A, 3L, function(x) solve(x), simplify = TRUE),
+           dim = rev(dim(lxd)))
> chk(a, b, check.attributes = FALSE)
> chk(solve(lxn, y), Mult(solve(lxn), y))
> chk(solve(lxd, y), Mult(solve(lxd), y))
> ### recycle C
> chk(solve(lxn[1,], y), as.array(solve(lxn[1,]))[,1] %% y)
> chk(solve(lxn[rep(1, N),], y), solve(lxn[1,], y), check.attributes = FALSE)
> ### recycle y
> chk(solve(lxn, y[,1]), solve(lxn, y[,rep(1, N)]))

    also for  $\mathbf{C}_i^\top \mathbf{x}_i = \mathbf{y}_i$ 

> chk(solve(lxn[1,], y, transpose = TRUE),
+     t(as.array(solve(lxn[1,]))[,1] %% y)

```

## 2.8 Log-determinants

For computing the log-determinant  $\log(\det(\mathbf{C}_i)) = \sum_{j=1}^J \log(\text{diag}(\mathbf{C}_i)_j)$  we sum over the log-diagonal entries of a lower triangular matrix in C, both when the data are stored in row- and column-major order:

$\langle \text{logdet } 33a \rangle \equiv$

```
SEXP R_ltMatrices_logdet (SEXP C, SEXP N, SEXP J, SEXP diag, SEXP byrow) {

  SEXP ans;
  double *dans;
  int i, j, k;

   $\langle RC \text{ input } 23b \rangle$ 
  Rboolean Rbyrow = asLogical(byrow);
   $\langle C \text{ length } 24a \rangle$ 

  PROTECT(ans = allocVector(REALSXP, iN));
  dans = REAL(ans);

  for (i = 0; i < iN; i++) {
    dans[i] = 0.0;
    if (Rdiag) {
      k = 1;
      for (j = 0; j < iJ; j++) {
        dans[i] += log(dC[k - 1]);
        k += (Rbyrow ? j + 2 : iJ - j);
      }
      dC += p;
    }
  }

  UNPROTECT(1);
  return(ans);
}
◇
```

Fragment referenced in 3.

The R interface now simply calls this low-level function

$\langle \text{logdet } \text{ltMatrices } 33b \rangle \equiv$

```
logdet <- function(x) {

  if (!is.ltMatrices(x))
    stop("x is not an ltMatrices object")

  byrow <- attr(x, "byrow")
  diag <- attr(x, "diag")
  d <- dim(x)
  J <- d[2L]
  dn <- dimnames(x)
  if (!is.double(x)) storage.mode(x) <- "double"

  ret <- .Call(mvtnorm_R_ltMatrices_logdet, x,
              as.integer(d[1L]), as.integer(J), as.logical(diag),
              as.logical(byrow))
  names(ret) <- dn[[1L]]
  return(ret)
}
◇
```

Fragment referenced in 2.

We test the functionality by extracting the diagonal elements from different matrices and summing over their logarithms

```
> chk(logdet(lxn), colSums(log(diagonals(lxn))))
> chk(logdet(lxd[1,]), colSums(log(diagonals(lxd[1,]))))
> chk(logdet(lxd), colSums(log(diagonals(lxd))))
> lxd2 <- ltMatrices(lxd, byrow = !attr(lxd, "byrow"))
> chk(logdet(lxd2), colSums(log(diagonals(lxd2))))
```

## 2.9 Crossproducts

We want to compute  $\mathbf{C}_i \mathbf{C}_i^\top$  or  $\text{diag}(\mathbf{C}_i \mathbf{C}_i^\top)$  (`diag_only = TRUE`) for  $i = 1, \dots, N$ . These are symmetric matrices, so we store them as a lower triangular matrix using a different class name `syMatrices`. We write one C function for computing  $\mathbf{C}_i \mathbf{C}_i^\top$  or  $\mathbf{C}_i^\top \mathbf{C}_i$  (`Rtranspose` being `TRUE`).

We differentiate between computation of the diagonal elements of the crossproduct

*<first element 34a>  $\equiv$*

```
    dans[0] = 1.0;
    if (Rdiag)
        dans[0] = pow(dC[0], 2);
    if (Rtranspose) { // crossprod
        for (k = 1; k < iJ; k++)
            dans[0] += pow(dC[IDX(k + 1, 1, iJ, Rdiag)], 2);
    }
    ◇
```

Fragment referenced in [34b](#), [35a](#).

*<tcrossprod diagonal only 34b>  $\equiv$*

```
    PROTECT(ans = allocMatrix(REALSXP, iJ, iN));
    dans = REAL(ans);
    for (n = 0; n < iN; n++) {
        <first element 34a>
        for (i = 1; i < iJ; i++) {
            dans[i] = 0.0;
            if (Rtranspose) { // crossprod
                for (k = i + 1; k < iJ; k++)
                    dans[i] += pow(dC[IDX(k + 1, i + 1, iJ, Rdiag)], 2);
            } else { // tcrossprod
                for (k = 0; k < i; k++)
                    dans[i] += pow(dC[IDX(i + 1, k + 1, iJ, Rdiag)], 2);
            }
            if (Rdiag) {
                dans[i] += pow(dC[IDX(i + 1, i + 1, iJ, Rdiag)], 2);
            } else {
                dans[i] += 1.0;
            }
        }
        dans += iJ;
        dC += len;
    }
    ◇
```

Fragment referenced in [36](#).

and computation of the full  $J \times J$  crossproduct matrix

$\langle \text{tcrossprod full 35a} \rangle \equiv$

```

nrow = iJ * (iJ + 1) / 2;
PROTECT(ans = allocMatrix(REALSXP, nrow, iN));
dans = REAL(ans);
for (n = 0; n < INTEGER(N)[0]; n++) {
   $\langle \text{first element 34a} \rangle$ 
  for (i = 1; i < iJ; i++) {
    for (j = 0; j <= i; j++) {
      ix = IDX(i + 1, j + 1, iJ, 1);
      dans[ix] = 0.0;
      if (Rtranspose) { // crossprod
        for (k = i + 1; k < iJ; k++)
          dans[ix] +=
            dC[IDX(k + 1, i + 1, iJ, Rdiag)] *
            dC[IDX(k + 1, j + 1, iJ, Rdiag)];
      } else { // tcrossprod
        for (k = 0; k < j; k++)
          dans[ix] +=
            dC[IDX(i + 1, k + 1, iJ, Rdiag)] *
            dC[IDX(j + 1, k + 1, iJ, Rdiag)];
      }
      if (Rdiag) {
        if (Rtranspose) {
          dans[ix] +=
            dC[IDX(i + 1, i + 1, iJ, Rdiag)] *
            dC[IDX(i + 1, j + 1, iJ, Rdiag)];
        } else {
          dans[ix] +=
            dC[IDX(i + 1, j + 1, iJ, Rdiag)] *
            dC[IDX(j + 1, j + 1, iJ, Rdiag)];
        }
      } else {
        if (j < i)
          dans[ix] += dC[IDX(i + 1, j + 1, iJ, Rdiag)];
        else
          dans[ix] += 1.0;
      }
    }
  }
  dans += nrow;
  dC += len;
}

```

◇

Fragment referenced in [36](#).

and put both cases together

$\langle \text{IDX 35b} \rangle \equiv$

```

#define IDX(i, j, n, d) ((i) >= (j) ? (n) * ((j) - 1) - ((j) - 2) * ((j) - 1)/2 + (i) - (j) - (!d) * (

```

◇

Fragment referenced in [36](#), [42a](#).

$\langle \text{tcrossprod } 36 \rangle \equiv$

$\langle \text{IDX } 35b \rangle$

```
SEXP R_ltMatrices_tcrossprod (SEXP C, SEXP N, SEXP J, SEXP diag,  
                             SEXP diag_only, SEXP transpose) {
```

```
    SEXP ans;  
    double *dans;  
    int i, j, n, k, ix, nrow;
```

$\langle \text{RC input } 23b \rangle$

```
Rboolean Rdiag_only = asLogical(diag_only);  
Rboolean Rtranspose = asLogical(transpose);
```

```
if (Rdiag_only) {  
     $\langle \text{tcrossprod diagonal only } 34b \rangle$   
} else {  
     $\langle \text{tcrossprod full } 35a \rangle$   
}  
UNPROTECT(1);  
return(ans);
```

```
}  
◇
```

Fragment referenced in 3.

with R interface

$\langle$  *tcrossprod* *ltMatrices* 37  $\rangle \equiv$

```
### C %*% t(C) => returns object of class syMatrices
### diag(C %*% t(C)) => returns matrix of diagonal elements
.Tcrossprod <- function(x, diag_only = FALSE, transpose = FALSE) {

  if (!is.ltMatrices(x)) {
    ret <- tcrossprod(x)
    if (diag_only) ret <- diag(ret)
    return(ret)
  }

  byrow_orig <- attr(x, "byrow")
  diag <- attr(x, "diag")
  d <- dim(x)
  N <- d[1L]
  J <- d[2L]
  dn <- dimnames(x)

  x <- ltMatrices(x, byrow = FALSE)
  if (!is.double(x)) storage.mode(x) <- "double"

  ret <- .Call(mvtnorm_R_ltMatrices_tcrossprod, x, as.integer(N), as.integer(J),
              as.logical(diag), as.logical(diag_only), as.logical(transpose))
  colnames(ret) <- dn[[1L]]
  if (diag_only) {
    rownames(ret) <- dn[[2L]]
  } else {
    ret <- ltMatrices(ret, diag = TRUE, byrow = FALSE, names = dn[[2L]])
    ret <- as.syMatrices(ltMatrices(ret, byrow = byrow_orig))
  }
  return(ret)
}
.Tcrossprod <- function(x, diag_only = FALSE)
  .Tcrossprod(x = x, diag_only = diag_only, transpose = FALSE)
◇
```

Fragment referenced in [2](#).

We could have created yet another generic `tcrossprod`, but `base::tcrossprod` is more general and, because speed is an issue, we don't want to waste time on methods dispatch.

```
> ## Tcrossprod
> a <- as.array(Tcrossprod(lxn))
> b <- array(apply(as.array(lxn), 3L, function(x) tcrossprod(x), simplify = TRUE),
+           dim = rev(dim(lxn)))
> chk(a, b, check.attributes = FALSE)
> # diagonal elements only
> d <- Tcrossprod(lxn, diag_only = TRUE)
> chk(d, apply(a, 3, diag))
> chk(d, diagonals(Tcrossprod(lxn)))
> a <- as.array(Tcrossprod(lxd))
> b <- array(apply(as.array(lxd), 3L, function(x) tcrossprod(x), simplify = TRUE),
+           dim = rev(dim(lxd)))
> chk(a, b, check.attributes = FALSE)
> # diagonal elements only
> d <- Tcrossprod(lxd, diag_only = TRUE)
```

```
> chk(d, apply(a, 3, diag))
> chk(d, diagonals(Tcrossprod(lxd)))
```

We also add `Crossprod`, which is a call to `Tcrossprod` with the `transpose` switch turned on

$\langle \text{crossprod } \text{ltMatrices } 38 \rangle \equiv$

```
Crossprod <- function(x, diag_only = FALSE)
  .Tcrossprod(x, diag_only = diag_only, transpose = TRUE)
```

◇

Fragment referenced in [2](#).

and run some checks

```
> ## Crossprod
> a <- as.array(Crossprod(lxn))
> b <- array(apply(as.array(lxn), 3L, function(x) crossprod(x), simplify = TRUE),
+           dim = rev(dim(lxn)))
> chk(a, b, check.attributes = FALSE)
> # diagonal elements only
> d <- Crossprod(lxn, diag_only = TRUE)
> chk(d, apply(a, 3, diag))
> chk(d, diagonals(Crossprod(lxn)))
> a <- as.array(Crossprod(lxd))
> b <- array(apply(as.array(lxd), 3L, function(x) crossprod(x), simplify = TRUE),
+           dim = rev(dim(lxd)))
> chk(a, b, check.attributes = FALSE)
> # diagonal elements only
> d <- Crossprod(lxd, diag_only = TRUE)
> chk(d, apply(a, 3, diag))
> chk(d, diagonals(Crossprod(lxd)))
```

## 2.10 Cholesky Factorisation

One might want to compute the Cholesky factorisations  $\Sigma_i = \mathbf{C}_i \mathbf{C}_i^\top$  for multiple symmetric matrices  $\Sigma_i$ , stored as a matrix in class `syMatrices`.

$\langle chol\ syMatrices\ 39 \rangle \equiv$

```
chol.syMatrices <- function(x, ...) {

  byrow_orig <- attr(x, "byrow")
  dnm <- dimnames(x)
  stopifnot(attr(x, "diag"))
  d <- dim(x)

  ### x is of class syMatrices, coerce to ltMatrices first and re-arrange
  ### second
  x <- ltMatrices(unclass(x), diag = TRUE,
                  byrow = byrow_orig, names = dnm[[2L]])
  x <- ltMatrices(x, byrow = FALSE)
  # class(x) <- class(x)[-1]
  if (!is.double(x)) storage.mode(x) <- "double"

  ret <- .Call(mvtnorm_R_syMatrices_chol, x,
              as.integer(d[1L]), as.integer(d[2L]))
  colnames(ret) <- dnm[[1L]]

  ret <- ltMatrices(ret, diag = TRUE,
                   byrow = FALSE, names = dnm[[2L]])
  ret <- ltMatrices(ret, byrow = byrow_orig)

  return(ret)
}
◇
```

Fragment referenced in [2](#).

Luckily, we already have the data in the correct packed column-major storage, so we swiftly loop over  $i = 1, \dots, N$  in C and hand over to LAPACK

$\langle chol\ 40 \rangle \equiv$

```
SEXP R_syMatrices_chol (SEXP Sigma, SEXP N, SEXP J) {

    SEXP ans;
    double *dans, *dSigma;
    int iJ = INTEGER(J)[0];
    int pJ = iJ * (iJ + 1) / 2;
    int iN = INTEGER(N)[0];
    int i, j, info = 0;
    char lo = 'L';

    PROTECT(ans = allocMatrix(REALSXP, pJ, iN));
    dans = REAL(ans);
    dSigma = REAL(Sigma);

    for (i = 0; i < iN; i++) {

        /* copy data */
        for (j = 0; j < pJ; j++)
            dans[j] = dSigma[j];

        F77_CALL(dpptrf)(&lo, &iJ, dans, &info FCONE);

        if (info != 0) {
            if (info > 0)
                error("the leading minor of order %d is not positive definite",
                    info);
            error("argument %d of Lapack routine %s had invalid value",
                -info, "dpptrf");
        }

        dSigma += pJ;
        dans += pJ;
    }
    UNPROTECT(1);
    return(ans);
}
◇
```

Fragment referenced in 3.

This new `chol` method can be used to revert `Tcrossprod` for `ltMatrices` with and without unit diagonals:

```
> Sigma <- Tcrossprod(lxd)
> chk(chol(Sigma), lxd)
> Sigma <- Tcrossprod(lxn)
> ## Sigma and chol(Sigma) always have diagonal, lxn doesn't
> chk(as.array(chol(Sigma)), as.array(lxn))
```

## 2.11 Kronecker Products

We sometimes need to compute  $\text{vec}(\mathbf{S})^\top (\mathbf{A}^\top \otimes \mathbf{C})$ , where  $\mathbf{S}$  is a lower triangular or other  $J \times J$  matrix and  $\mathbf{A}$  and  $\mathbf{C}$  are lower triangular  $J \times J$  matrices. With the “vec trick”, we have  $\text{vec}(\mathbf{S})^\top (\mathbf{A}^\top \otimes \mathbf{C}) = \text{vec}(\mathbf{C}^\top \mathbf{S} \mathbf{A}^\top)^\top$ . The LAPACK function `dtrmm` computes products of lower triangular matrices with other matrices, so we simply call this function looping over  $i = 1, \dots, N$ .

$\langle t(C) S t(A) 41 \rangle \equiv$

```

char siR = 'R', siL = 'L', lo = 'L', tr = 'N', trT = 'T', di = 'N', trs;
double ONE = 1.0;
int iJ2 = iJ * iJ;

double tmp[iJ2];
for (j = 0; j < iJ2; j++) tmp[j] = 0.0;

ans = PROTECT(allocMatrix(REALSXP, iJ2, iN));
dans = REAL(ans);

for (i = 0; i < LENGTH(ans); i++) dans[i] = 0.0;

for (i = 0; i < iN; i++) {

    /* A := C */
    for (j = 0; j < iJ; j++) {
        for (k = 0; k <= j; k++)
            tmp[k * iJ + j] = dC[IDX(j + 1, k + 1, iJ, 1L)];
    }

    /* S was already expanded in R code; B = S */
    for (j = 0; j < iJ2; j++) dans[j] = dS[j];

    /* B := t(A) %*% B */
    trs = (RtC ? trT : tr);
    F77_CALL(dtrmm)(&siL, &lo, &trs, &di, &iJ, &iJ, &ONE, tmp, &iJ,
                    dans, &iJ FCONE FCONE FCONE FCONE);

    /* A */
    for (j = 0; j < iJ; j++) {
        for (k = 0; k <= j; k++)
            tmp[k * iJ + j] = dA[IDX(j + 1, k + 1, iJ, 1L)];
    }

    /* B := B %*% t(A) */
    trs = (RtA ? trT : tr);
    F77_CALL(dtrmm)(&siR, &lo, &trs, &di, &iJ, &iJ, &ONE, tmp, &iJ,
                    dans, &iJ FCONE FCONE FCONE FCONE);

    dans += iJ2;
    dC += p;
    dS += iJ2;
    dA += p;
}
◇

```

Fragment referenced in [42a](#).

$\langle \text{vec trick 42a} \rangle \equiv$

$\langle \text{IDX 35b} \rangle$

```
SEXP R_vectrick(SEXP C, SEXP N, SEXP J, SEXP S, SEXP A, SEXP diag, SEXP trans) {

    int i, j, k;
    SEXP ans;
    double *dS, *dans, *dA;

    /* note: diag is needed by this chunk but has no consequences */
     $\langle RC \text{ input 23b} \rangle$ 
     $\langle C \text{ length 24a} \rangle$ 
    dS = REAL(S);
    dA = REAL(A);

    Rboolean RtC = LOGICAL(trans)[0];
    Rboolean RtA = LOGICAL(trans)[1];

     $\langle t(C) S t(A) 41 \rangle$ 

    UNPROTECT(1);
    return(ans);
}
◇
```

Fragment referenced in 3.

In R, we compute  $\mathbf{C}^\top \mathbf{S} \mathbf{A}^\top$  by default or  $\mathbf{C} \mathbf{S} \mathbf{A}^\top$  or  $\mathbf{C}^\top \mathbf{S} \mathbf{A}$  or  $\mathbf{C}^\top \mathbf{S} \mathbf{A}^\top$  by using the `trans` argument in `vectrick`. Argument `C` is an `ltMatrices` object

$\langle \text{check } C \text{ argument 42b} \rangle \equiv$

```
C <- as.ltMatrices(C)
if (!attr(C, "diag")) diagonals(C) <- 1
C_byrow_orig <- attr(C, "byrow")
C <- ltMatrices(C, byrow = FALSE)
dC <- dim(C)
nm <- attr(C, "rcnames")
N <- dC[1L]
J <- dC[2L]
class(C) <- class(C)[-1L] ### works because of as.ltMatrices(c)
if (!is.double(C)) storage.mode(C) <- "double"
◇
```

Fragment referenced in 44.

`S` can be an `ltMatrices` object or a  $J^2 \times N$  matrix featuring columns of vectorised  $J \times J$  matrices

*< check S argument 43a > ≡*

```

SltM <- is.ltMatrices(S)
if (SltM) {
  if (!attr(S, "diag")) diagonals(S) <- 1
  S_byrow_orig <- attr(S, "byrow")
  stopifnot(S_byrow_orig == C_byrow_orig)
  S <- ltMatrices(S, byrow = FALSE)
  dS <- dim(S)
  stopifnot(dC[2L] == dS[2L])
  if (dC[1L] != 1L) {
    stopifnot(dC[1L] == dS[1L])
  } else {
    N <- dS[1L]
  }
  ## argument A in dtrmm is not in packed form, so expand in J x J
  ## matrix
  S <- matrix(as.array(S), ncol = dS[1L])
} else {
  stopifnot(is.matrix(S))
  stopifnot(nrow(S) == J^2)
  if (dC[1L] != 1L) {
    stopifnot(dC[1L] == ncol(S))
  } else {
    N <- ncol(S)
  }
}
if (!is.double(S)) storage.mode(S) <- "double"
◇

```

Fragment referenced in [44](#).

A is an ltMatrices object

*< check A argument 43b > ≡*

```

if (missing(A)) {
  A <- C
} else {
  A <- as.ltMatrices(A)
  if (!attr(A, "diag")) diagonals(A) <- 1
  A_byrow_orig <- attr(A, "byrow")
  stopifnot(C_byrow_orig == A_byrow_orig)
  A <- ltMatrices(A, byrow = FALSE)
  dA <- dim(A)
  stopifnot(dC[2L] == dA[2L])
  class(A) <- class(A)[-1L]
  if (!is.double(A)) storage.mode(A) <- "double"
  if (dC[1L] != dA[1L]) {
    if (dC[1L] == 1L)
      C <- C[, rep(1, N), drop = FALSE]
    if (dA[1L] == 1L)
      A <- A[, rep(1, N), drop = FALSE]
    stopifnot(ncol(A) == ncol(C))
  }
}
◇

```

Fragment referenced in [44](#).

We put everything together in function `vecrick`

$\langle \text{kronecker vec trick 44} \rangle \equiv$

```
vecrick <- function(C, S, A, transpose = c(TRUE, TRUE)) {

  stopifnot(all(is.logical(transpose)))
  stopifnot(length(transpose) == 2L)

   $\langle \text{check } C \text{ argument 42b} \rangle$ 
   $\langle \text{check } S \text{ argument 43a} \rangle$ 
   $\langle \text{check } A \text{ argument 43b} \rangle$ 

  ret <- .Call(mvtnorm_R_vecrick, C, as.integer(N), as.integer(J), S, A,
              as.logical(TRUE), as.logical(transpose))

  if (!SltM) return(matrix(c(ret), ncol = N))

  L <- matrix(1:(J^2), nrow = J)
  ret <- ltMatrices(ret[L[lower.tri(L, diag = TRUE)],,drop = FALSE],
                  diag = TRUE, byrow = FALSE, names = nm)
  ret <- ltMatrices(ret, byrow = C_byrow_orig)
  return(ret)
}
◇
```

Fragment referenced in [2](#).

Here is a small example

```
> J <- 10
> d <- TRUE
> L <- diag(J)
> L[lower.tri(L, diag = d)] <- prm <- runif(J * (J + c(-1, 1)[d + 1]) / 2)
> C <- solve(L)
> D <- -kronecker(t(C), C)
> S <- diag(J)
> S[lower.tri(S, diag = TRUE)] <- x <- runif(J * (J + 1) / 2)
> SD0 <- matrix(c(S) %*% D, ncol = J)
> SD1 <- -crossprod(C, tcrossprod(S, C))
> a <- ltMatrices(C[lower.tri(C, diag = TRUE)], diag = TRUE, byrow = FALSE)
> b <- ltMatrices(x, diag = TRUE, byrow = FALSE)
> SD2 <- -vecrick(a, b, a)
> SD2a <- -vecrick(a, b)
> chk(SD2, SD2a)
> chk(SD0[lower.tri(SD0, diag = d)],
+     SD1[lower.tri(SD1, diag = d)])
> chk(SD0[lower.tri(SD0, diag = d)],
+     c(unclass(SD2)))
> ### same; but SD2 is vec(SD0)
> S <- t(matrix(as.array(b), byrow = FALSE, nrow = 1))
> SD2 <- -vecrick(a, S, a)
> SD2a <- -vecrick(a, S)
> chk(SD2, SD2a)
> chk(c(SD0), c(SD2))
> ### N > 1
```

```

> N <- 4L
> prm <- runif(J * (J - 1) / 2)
> C <- ltMatrices(prm)
> S <- matrix(runif(J^2 * N), ncol = N)
> A <- vectrick(C, S, C)
> Cx <- as.array(C)[,,1]
> B <- apply(S, 2, function(x) t(Cx) %% matrix(x, ncol = J) %% t(Cx))
> chk(A, B)
> A <- vectrick(C, S, C, transpose = c(FALSE, FALSE))
> Cx <- as.array(C)[,,1]
> B <- apply(S, 2, function(x) Cx %% matrix(x, ncol = J) %% Cx)
> chk(A, B)

```

## 2.12 Convenience Functions

We add a few convenience functions for computing covariance matrices  $\Sigma_i = \mathbf{C}_i \mathbf{C}_i^\top$ , precision matrices  $\mathbf{P}_i = \mathbf{L}_i^\top \mathbf{L}_i$ , correlation matrices  $\mathbf{R}_i = \tilde{\mathbf{C}}_i \tilde{\mathbf{C}}_i^\top$  (where  $\tilde{\mathbf{C}}_i = \text{diag}(\mathbf{C}_i \mathbf{C}_i^\top)^{-\frac{1}{2}} \mathbf{C}_i$ ), or matrices of partial correlations  $\mathbf{A}_i = -\tilde{\mathbf{L}}_i^\top \tilde{\mathbf{L}}_i$  with  $\tilde{\mathbf{L}}_i = \mathbf{L}_i \text{diag}(\mathbf{L}_i^\top \mathbf{L}_i)^{-\frac{1}{2}}$  from  $\mathbf{L}_i$  (`invchol`) or  $\mathbf{C}_i = \mathbf{L}_i^{-1}$  (`chol`) for  $i = 1, \dots, N$ .

Before we start, let us put a label on lower triangular matrices, such that we can differentiate between  $\mathbf{C}$  and  $\mathbf{L}$ .

$\langle \text{chol classes 45} \rangle \equiv$

```

is.chol <- function(x) inherits(x, "chol")
as.chol <- function(x) {
  stopifnot(is.ltMatrices(x))
  if (is.chol(x)) return(x)
  if (is.invchol(x))
    return(invchol2chol(x))
  class(x) <- c("chol", class(x))
  return(x)
}
is.invchol <- function(x) inherits(x, "invchol")
as.invchol <- function(x) {
  stopifnot(is.ltMatrices(x))
  if (is.invchol(x)) return(x)
  if (is.chol(x))
    return(chol2invchol(x))
  class(x) <- c("invchol", class(x))
  return(x)
}

```

Fragment referenced in [48](#).

First, we set-up functions for computing  $\tilde{\mathbf{C}}_i$

$\langle D \text{ times } C \text{ 46} \rangle \equiv$

```
Dchol <- function(x, D = 1 / sqrt(Tcrossprod(x, diag_only = TRUE))) {
  if (is.invchol(x)) stop("Dchol cannot work with invchol objects")
  x <- .adddiag(x)
  byrow_orig <- attr(x, "byrow")
  x <- ltMatrices(x, byrow = TRUE)
  N <- dim(x)[1L]
  J <- dim(x)[2L]
  nm <- dimnames(x)[[2L]]
  ### for some parameter configurations logdet(ret) would
  ### be -Inf; make sure this doesn't happen
  if (any(D < .Machine$double.eps))
    D[D < .Machine$double.eps] <- 2 * .Machine$double.eps
  x <- unclass(x) * D[rep(1:J, 1:J),,drop = FALSE]
  ret <- ltMatrices(x, diag = TRUE, byrow = TRUE, names = nm)
  ret <- as.chol(ltMatrices(ret, byrow = byrow_orig))
  return(ret)
}
◇
```

Fragment referenced in [48](#).

and  $\tilde{\mathbf{C}}_i^{-1} = \mathbf{L}_i \text{diag}(\mathbf{L}_i^{-1} \mathbf{L}_i^{-\top})^{\frac{1}{2}}$

$\langle L \text{ times } D \text{ 47} \rangle \equiv$

```
### invcholD = solve(Dchol)
invcholD <- function(x, D = sqrt(Tcrossprod(solve(x), diag_only = TRUE))) {

  if (is.chol(x)) stop("invcholD cannot work with chol objects")

  x <- .adddiag(x)

  byrow_orig <- attr(x, "byrow")

  x <- ltMatrices(x, byrow = FALSE)

  N <- dim(x)[1L]
  J <- dim(x)[2L]
  nm <- dimnames(x)[[2L]]

  ### for some parameter configurations logdet(ret) would
  ### be -Inf; make sure this doesn't happen
  if (any(D < .Machine$double.eps))
    D[D < .Machine$double.eps] <- 2 * .Machine$double.eps

  x <- unclass(x) * D[rep(1:J, J:1),,drop = FALSE]

  ret <- ltMatrices(x, diag = TRUE, byrow = FALSE, names = nm)
  ret <- as.invchol(ltMatrices(ret, byrow = byrow_orig))
  return(ret)
}
◇
```

Fragment referenced in [48](#).

and now the convenience functions are one-liners:

$\langle$  *convenience functions* 48  $\rangle \equiv$

```
 $\langle$  chol classes 45  $\rangle$ 
 $\langle$  D times C 46  $\rangle$ 
 $\langle$  L times D 47  $\rangle$ 

### C -> Sigma
chol2cov <- function(x)
  Tcrossprod(x)

### L -> C
invchol2chol <- function(x)
  as.chol(solve(x))

### C -> L
chol2invchol <- function(x)
  as.invchol(solve(x))

### L -> Sigma
invchol2cov <- function(x)
  chol2cov(invchol2chol(x))

### L -> Precision
invchol2pre <- function(x)
  Crossprod(x)

### C -> Precision
chol2pre <- function(x)
  Crossprod(chol2invchol(x))

### C -> R
chol2cor <- function(x) {
  ret <- Tcrossprod(Dchol(x))
  diagonals(ret) <- NULL
  return(ret)
}

### L -> R
invchol2cor <- function(x) {
  ret <- chol2cor(invchol2chol(x))
  diagonals(ret) <- NULL
  return(ret)
}

### L -> A
invchol2pc <- function(x) {
  ret <- -Crossprod(invcholD(x, D = 1 / sqrt(Crossprod(x, diag_only = TRUE))))
  diagonals(ret) <- 0
  ret
}

### C -> A
chol2pc <- function(x)
  invchol2pc(solve(x))
 $\diamond$ 
```

Fragment referenced in 2.

Here are some tests

```
> prec2pc <- function(x) {
+   ret <- -cov2cor(x)
+   diag(ret) <- 0
+   ret
+ }
> L <- lxn
> Sigma <- apply(as.array(L), 3,
+   function(x) tcrossprod(solve(x)), simplify = FALSE)
> Prec <- lapply(Sigma, solve)
> Corr <- lapply(Sigma, cov2cor)
> CP <- lapply(Corr, solve)
> PC <- lapply(Prec, function(x) prec2pc(x))
> chk(unlist(Sigma), c(as.array(invchol2cov(L))),
+   check.attributes = FALSE)
> chk(unlist(Prec), c(as.array(invchol2pre(L))),
+   check.attributes = FALSE)
> chk(unlist(Corr), c(as.array(invchol2cor(L))),
+   check.attributes = FALSE)
> chk(unlist(CP), c(as.array(Crossprod(invcholD(L)))),
+   check.attributes = FALSE)
> chk(unlist(PC), c(as.array(invchol2pc(L))),
+   check.attributes = FALSE)

> C <- lxn
> Sigma <- apply(as.array(C), 3,
+   function(x) tcrossprod(x), simplify = FALSE)
> Prec <- lapply(Sigma, solve)
> Corr <- lapply(Sigma, cov2cor)
> CP <- lapply(Corr, solve)
> PC <- lapply(Prec, function(x) prec2pc(x))
> chk(unlist(Sigma), c(as.array(chol2cov(C))),
+   check.attributes = FALSE)
> chk(unlist(Prec), c(as.array(chol2pre(C))),
+   check.attributes = FALSE)
> chk(unlist(Corr), c(as.array(chol2cor(C))),
+   check.attributes = FALSE)
> chk(unlist(CP), c(as.array(Crossprod(solve(Dchol(C))))),
+   check.attributes = FALSE)
> chk(unlist(PC), c(as.array(chol2pc(C))),
+   check.attributes = FALSE)

> L <- lxd
> Sigma <- apply(as.array(L), 3,
+   function(x) tcrossprod(solve(x)), simplify = FALSE)
> Prec <- lapply(Sigma, solve)
> Corr <- lapply(Sigma, cov2cor)
> CP <- lapply(Corr, solve)
> PC <- lapply(Prec, function(x) prec2pc(x))
> chk(unlist(Sigma), c(as.array(invchol2cov(L))),
+   check.attributes = FALSE)
> chk(unlist(Prec), c(as.array(invchol2pre(L))),
+   check.attributes = FALSE)
> chk(unlist(Corr), c(as.array(invchol2cor(L))),
```

```

+     check.attributes = FALSE)
> chk(unlist(CP), c(as.array(Crossprod(invcholD(L)))),
+     check.attributes = FALSE)
> chk(unlist(PC), c(as.array(invchol2pc(L))),
+     check.attributes = FALSE)

> C <- lxd
> Sigma <- apply(as.array(C), 3,
+     function(x) tcrossprod(x), simplify = FALSE)
> Prec <- lapply(Sigma, solve)
> Corr <- lapply(Sigma, cov2cor)
> CP <- lapply(Corr, solve)
> PC <- lapply(Prec, function(x) prec2pc(x))
> chk(unlist(Sigma), c(as.array(chol2cov(C))),
+     check.attributes = FALSE)
> chk(unlist(Prec), c(as.array(chol2pre(C))),
+     check.attributes = FALSE)
> chk(unlist(Corr), c(as.array(chol2cor(C))),
+     check.attributes = FALSE)
> chk(unlist(CP), c(as.array(Crossprod(solve(Dchol(C))))),
+     check.attributes = FALSE)
> chk(unlist(PC), c(as.array(chol2pc(C))),
+     check.attributes = FALSE)

```

We also add an `aperm` method for class `ltMatrices`, implementing the parameters  $(\tilde{\mathbf{C}}_i$  or  $\tilde{\mathbf{L}}_i)$  for permuted versions of the random vectors  $\mathbf{Y}_i$ . Let  $\pi$  denote a permutation of  $1, \dots, J$  and  $\Pi$  the corresponding permutation matrix. Then, we have  $\Pi \mathbf{Y}_i \sim \mathbb{N}_J(\mathbf{0}_J, \Pi \mathbf{C}_i \mathbf{C}_i^\top \Pi^\top)$ . Unfortunately,  $\Pi \mathbf{C}_i$  is no longer lower triangular, so we have to find the Cholesky decomposition  $\tilde{\mathbf{C}}_i \tilde{\mathbf{C}}_i^\top$  of  $\Pi \mathbf{C}_i \mathbf{C}_i^\top \Pi^\top$ . Of course,  $\tilde{\mathbf{L}}_i = \tilde{\mathbf{C}}_i^{-1}$ .

The function `aperm`, with argument `perm =  $\pi$` , now computes the Cholesky factor  $\tilde{\mathbf{C}}_i$  of the permuted covariance matrix, or the inverse thereof (in case `x` is of class `invchol`). We start with some tests

`< aperm checks 50 >  $\equiv$`

```

J <- dim(a)[2L]
if (missing(perm)) return(a)
if (is.character(perm))
  perm <- match(perm, dimnames(a)[[2L]])
stopifnot(all(perm %in% 1:J))

args <- list(...)
if (length(args) > 0L)
  warning("Additional arguments", names(args), "ignored")

```

Fragment referenced in [51a](#).

and then implement the two methods

$\langle \text{aperm 51a} \rangle \equiv$

```
aperm.chol <- function(a, perm, ...) {
   $\langle \text{aperm checks 50} \rangle$ 
  return(as.chol(chol(chol2cov(a)[,perm])))
}
aperm.invchol <- function(a, perm, ...) {
   $\langle \text{aperm checks 50} \rangle$ 
  return(chol2invchol(chol(invchol2cov(a)[,perm])))
}
 $\diamond$ 
```

Fragment defined by [51ab](#).

Fragment referenced in [2](#).

```
> L <- as.invchol(1xn)
> J <- dim(L)[2L]
> Lp <- aperm(a = L, perm = p <- sample(1:J))
> chk(invchol2cov(L)[,p], invchol2cov(Lp))
> C <- as.chol(1xn)
> J <- dim(C)[2L]
> Cp <- aperm(a = C, perm = p <- sample(1:J))
> chk(chol2cov(C)[,p], chol2cov(Cp))
```

We finally add a method for class `ltMatrices`, for which we actually cannot provide a reasonable result, and for symmetric matrices, where we simply fall-back on subsetting

$\langle \text{aperm 51b} \rangle \equiv$

```
aperm.ltMatrices <- function(a, perm, ...)
  stop("Cannot permute objects of class ltMatrices,
       consider calling as.chol() or as.invchol() first")

aperm.syMatrices <- function(a, perm, ...)
  return(a[,perm])
 $\diamond$ 
```

Fragment defined by [51ab](#).

Fragment referenced in [2](#).

## 2.13 Marginal and Conditional Normal Distributions

Marginal and conditional distributions from distributions  $\mathbf{Y}_i \sim \mathbb{N}_J(\mathbf{0}_J, \mathbf{C}_i \mathbf{C}_i^\top)$  (`chol` argument for  $\mathbf{C}_i$  for  $i = 1, \dots, N$ ) or  $\mathbf{Y}_i \sim \mathbb{N}_J(\mathbf{0}_J, \mathbf{L}_i^{-1} \mathbf{L}_i^{-\top})$  (`invchol` argument for  $\mathbf{L}_i$  for  $i = 1, \dots, N$ ) shall be computed.

$\langle mc \text{ input checks } 52a \rangle \equiv$

```
stopifnot(xor(missing(chol), missing(invchol)))
x <- if (missing(chol)) invchol else chol

stopifnot(is.ltMatrices(x))

N <- dim(x)[1L]
J <- dim(x)[2L]

if (missing(which)) return(x)

if (is.character(which)) which <- match(which, dimnames(x)[[2L]])
stopifnot(all(which %in% 1:J))
◇
```

Fragment referenced in [52b](#), [55](#).

The first  $j$  marginal distributions can be obtained from subsetting  $\mathbf{C}$  or  $\mathbf{L}$  directly. Arbitrary marginal distributions are based on the corresponding subset of the covariance matrix for which we compute a corresponding Cholesky factor (such that we can use `lpmvnorm` later on).

$\langle marginal \ 52b \rangle \equiv$

```
marg_mvnorm <- function(chol, invchol, which = 1L) {

   $\langle mc \text{ input checks } 52a \rangle$ 

  if (which[1] == 1L && (length(which) == 1L ||
    all(diff(which) == 1L))) {
    ### which is 1:j
    tmp <- x[,which]
  } else {
    if (missing(chol)) x <- invchol2chol(x)
    ### note: aperm would work but computes
    ### Cholesky of J^2, here only length(which)^2
    ### is needed
    tmp <- base::chol(chol2cov(x)[,which])
    if (missing(chol)) tmp <- chol2invchol(tmp)
  }

  if (missing(chol))
    ret <- list(invchol = tmp)
  else
    ret <- list(chol = tmp)

  ret
}
◇
```

Fragment referenced in [2](#).

We compute conditional distributions from the precision matrices  $\Sigma_i^{-1} = \mathbf{P}_i = \mathbf{L}_i^\top \mathbf{L}_i$  (we omit the  $i$  index from now on). For an arbitrary subset  $\mathbf{j} \subset \{1, \dots, J\}$ , the conditional distribution of  $\mathbf{Y}_{-\mathbf{j}}$  given  $\mathbf{Y}_{\mathbf{j}} = \mathbf{y}_{\mathbf{j}}$  is

$$\mathbf{Y}_{-\mathbf{j}} \mid \mathbf{Y}_{\mathbf{j}} = \mathbf{y}_{\mathbf{j}} \sim \mathbb{N}_{|\mathbf{j}|} \left( -\mathbf{P}_{-\mathbf{j},-\mathbf{j}}^{-1} \mathbf{P}_{-\mathbf{j},\mathbf{j}} \mathbf{y}_{\mathbf{j}}, \mathbf{P}_{-\mathbf{j},-\mathbf{j}}^{-1} \right)$$

and we return a Cholesky factor  $\tilde{\mathbf{C}}$  such that  $\mathbf{P}_{-j,-j}^{-1} = \tilde{\mathbf{C}}\tilde{\mathbf{C}}^\top$  (if `chol` was given) or  $\tilde{\mathbf{L}} = \tilde{\mathbf{C}}^{-1}$  (if `invchol` was given).

We can implement this as

*(cond general 53)*  $\equiv$

```
stopifnot(!center)

if (!missing(chol)) ### chol is C = Cholesky of covariance
  P <- Crossprod(solve(chol)) ### P = t(L) %*% L with L = C^-1
else
  ### invchol is L = Cholesky of precision
  P <- Crossprod(invchol)

Pw <- P[, -which]
chol <- solve(base::chol(Pw))
Pa <- as.array(P)
Sa <- as.array(S <- Crossprod(chol))
if (dim(chol)[1L] == 1L) {
  Pa <- Pa[, , 1]
  Sa <- Sa[, , 1]
  mean <- -Sa %*% Pa[-which, which, drop = FALSE] %*% given
} else {
  if (ncol(given) == N) {
    mean <- sapply(1:N, function(i)
      -Sa[, , i] %*% Pa[-which, which, i] %*% given[, i, drop = FALSE])
  } else { ### compare to Mult() with ncol(y) !in% (1, N)
    mean <- sapply(1:N, function(i)
      -Sa[, , i] %*% Pa[-which, which, i] %*% given)
  }
}
}
```

Fragment referenced in 55.

If  $\mathbf{j} = \{1, \dots, j < J\}$  and  $\mathbf{L}$  is given, computations simplify a lot because the conditional precision matrix is

$$\mathbf{P}_{-j,-j} = (\mathbf{L}^\top \mathbf{L})_{-j,-j} = \mathbf{L}_{-j,-j}^\top \mathbf{L}_{-j,-j}$$

and thus we return  $\tilde{\mathbf{L}} = \mathbf{L}_{-j,-j}$  (if `invchol` was given) or  $\tilde{\mathbf{C}} = \mathbf{L}_{-j,-j}^{-1}$  (if `chol` was given). The conditional mean is

$$\begin{aligned} -\mathbf{P}_{-j,-j}^{-1} \mathbf{P}_{-j,j} \mathbf{y}_j &= -\mathbf{L}_{-j,-j}^{-1} \mathbf{L}_{-j,-j}^{-\top} \mathbf{L}_{-j,-j}^\top \mathbf{L}_{-j,j} \mathbf{y}_j \\ &= -\mathbf{L}_{-j,-j}^{-1} \mathbf{L}_{-j,j} \mathbf{y}_j. \end{aligned}$$

We sometimes, for example when scores with respect to  $\mathbf{L}_{-j,-j}^{-1}$  shall be computed in `slpmvnorm`, need the negative rescaled mean  $\mathbf{L}_{-j,j} \mathbf{y}_j$  and the `center = TRUE` argument triggers this values to be returned.

The implementation reads

$\langle \text{cond simple 54} \rangle \equiv$

```

if (which[1] == 1L && (length(which) == 1L ||
    all(diff(which) == 1L))) {
  ### which is 1:j
  L <- if (missing(invchol)) solve(chol) else invchol
  tmp <- matrix(0, ncol = ncol(given), nrow = J - length(which))
  centerm <- Mult(L, rbind(given, tmp))
  ### if ncol(given) is not N = dim(L)[1L] > 1, then
  ### solve() below won't work and we loop over
  ### columns of centerm
  if (dim(L)[1L] > 1 && ncol(given) != N) {
    centerm <- lapply(1:ncol(centerm), function(j)
      matrix(centerm[,j], nrow = J, ncol = N)[-which,,drop = FALSE]
    )
  } else {
    centerm <- centerm[-which,,drop = FALSE]
  }
  L <- L[,-which]
  ct <- centerm
  if (!is.matrix(ct)) ct <- do.call("rbind", ct)
  if (is.matrix(centerm)) {
    m <- -solve(L, centerm)
  } else {
    m <- do.call("rbind", lapply(centerm, function(cm) -solve(L, cm)))
  }
  if (missing(invchol)) {
    if (center)
      return(list(center = ct, chol = solve(L)))
    return(list(mean = m, chol = solve(L)))
  }
  if (center)
    return(list(center = ct, invchol = L))
  return(list(mean = m, invchol = L))
}

```

◇

Fragment referenced in [55](#).

Note that we could have avoided the general case altogether by first computing a Cholesky decomposition of the permuted covariance matrix (such that the conditioning variables come first). The code above only decomposes the marginal (and thus lower-dimensional) covariance. However, we didn't implement the `center = TRUE` case, so we can fall back on the permuted version if this option is requested. Putting everything together gives

$\langle \text{conditional } 55 \rangle \equiv$

```
cond_mvnorm <- function(chol, invchol, which_given = 1L, given, center = FALSE) {

  which <- which_given
   $\langle \text{mc input checks } 52a \rangle$ 

  if (N == 1) N <- NCOL(given)
  stopifnot(is.matrix(given) && nrow(given) == length(which))

   $\langle \text{cond simple } 54 \rangle$ 

  ### general with center = TRUE => permute first and go simple
  if (center) {
    perm <- c(which, (1:J)[!(1:J) %in% which])
    if (!missing(chol))
      return(cond_mvnorm(chol = aperm(as.chol(chol), perm = perm),
                           which_given = 1:length(which), given = given,
                           center = center))
    return(cond_mvnorm(invchol = aperm(as.invchol(invchol), perm = perm),
                        which_given = 1:length(which), given = given,
                        center = center))
  }

   $\langle \text{cond general } 53 \rangle$ 

  chol <- base::chol(S)
  if (missing(invchol))
    return(list(mean = mean, chol = chol))

  return(list(mean = mean, invchol = solve(chol)))
}
 $\diamond$ 
```

Fragment referenced in 2.

Let's check this against the commonly used formula based on the covariance matrix, first for the marginal distribution

```
> Sigma <- Tcrossprod(1xd)
> j <- 1:3
> chk(Sigma[,j], Tcrossprod(marg_mvnorm(chol = 1xd, which = j)$chol))
> j <- 2:4
> chk(Sigma[,j], Tcrossprod(marg_mvnorm(chol = 1xd, which = j)$chol))
> Sigma <- Tcrossprod(solve(1xd))
> j <- 1:3
> chk(Sigma[,j], Tcrossprod(solve(marg_mvnorm(invchol = 1xd, which = j)$invchol)))
> j <- 2:4
> chk(Sigma[,j], Tcrossprod(solve(marg_mvnorm(invchol = 1xd, which = j)$invchol)))
```

and then for conditional distributions. The general case is

```
> Sigma <- as.array(Tcrossprod(1xd))[, , 1]
> j <- 2:4
> y <- matrix(c(-1, 2, 1), nrow = 3)
> cm <- Sigma[-j, j, drop = FALSE] %*% solve(Sigma[j, j]) %*% y
> cS <- Sigma[-j, -j] - Sigma[-j, j, drop = FALSE] %*%
```

```

+       solve(Sigma[j,j]) %*% Sigma[j,-j,drop = FALSE]
> cmv <- cond_mvnorm(chol = lxd[1,], which = j, given = y)
> chk(cm, cmv$mean)
> chk(cS, as.array(Tcrossprod(cmv$chol))[,1])
> Sigma <- as.array(Tcrossprod(solve(lxd))[,1])
> j <- 2:4
> y <- matrix(c(-1, 2, 1), nrow = 3)
> cm <- Sigma[-j, j,drop = FALSE] %*% solve(Sigma[j,j]) %*% y
> cS <- Sigma[-j, -j] - Sigma[-j,j,drop = FALSE] %*%
+       solve(Sigma[j,j]) %*% Sigma[j,-j,drop = FALSE]
> cmv <- cond_mvnorm(invchol = lxd[1,], which = j, given = y)
> chk(cm, cmv$mean)
> chk(cS, as.array(Tcrossprod(solve(cmv$invchol))[,1])

```

and the simple case is

```

> Sigma <- as.array(Tcrossprod(lxd))[,1]
> j <- 1:3
> y <- matrix(c(-1, 2, 1), nrow = 3)
> cm <- Sigma[-j, j,drop = FALSE] %*% solve(Sigma[j,j]) %*% y
> cS <- Sigma[-j, -j] - Sigma[-j,j,drop = FALSE] %*%
+       solve(Sigma[j,j]) %*% Sigma[j,-j,drop = FALSE]
> cmv <- cond_mvnorm(chol = lxd[1,], which = j, given = y)
> chk(c(cm), c(cmv$mean))
> chk(cS, as.array(Tcrossprod(cmv$chol))[,1])
> Sigma <- as.array(Tcrossprod(solve(lxd))[,1])
> j <- 1:3
> y <- matrix(c(-1, 2, 1), nrow = 3)
> cm <- Sigma[-j, j,drop = FALSE] %*% solve(Sigma[j,j]) %*% y
> cS <- Sigma[-j, -j] - Sigma[-j,j,drop = FALSE] %*%
+       solve(Sigma[j,j]) %*% Sigma[j,-j,drop = FALSE]
> cmv <- cond_mvnorm(invchol = lxd[1,], which = j, given = y)
> chk(c(cm), c(cmv$mean))
> chk(cS, as.array(Tcrossprod(solve(cmv$invchol))[,1])

```

## 2.14 Continuous Log-likelihoods

With  $\mathbf{Z} \sim \mathbb{N}_J(0, \mathbf{I}_J)$  and  $\mathbf{Y} = \mathbf{C}_i \mathbf{Z} + \boldsymbol{\mu}_i \sim \mathbb{N}_J(\boldsymbol{\mu}_i, \mathbf{C}_i \mathbf{C}_i^\top)$  we want to evaluate the log-likelihood contributions for observations  $\mathbf{y}_1, \dots, \mathbf{y}_N$  in a function called `ldmvnorm`

$\langle \text{ldmvnorm 57a} \rangle \equiv$

```
ldmvnorm <- function(obs, mean = 0, chol, invchol, logLik = TRUE) {

  stopifnot(xor(missing(chol), missing(invchol)))
  if (!is.matrix(obs)) obs <- matrix(obs, ncol = 1L)
  p <- ncol(obs)

  if (!missing(chol)) {
     $\langle \text{ldmvnorm chol 59a} \rangle$ 
  } else {
     $\langle \text{ldmvnorm invchol 59b} \rangle$ 
  }

  names(logretval) <- colnames(obs)
  if (logLik) return(sum(logretval))
  return(logretval)
}
\diamond
```

Fragment referenced in [64](#).

We first check if the observations  $\mathbf{y}_1, \dots, \mathbf{y}_N$  are given in an  $J \times N$  matrix `obs` with corresponding means  $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_N$  in `means`.

$\langle \text{check obs 57b} \rangle \equiv$

```
.check_obs <- function(obs, mean, J, N) {

  nr <- nrow(obs)
  nc <- ncol(obs)
  if (nc != N)
    stop("obs and (inv)chol have non-conforming size")
  if (nr != J)
    stop("obs and (inv)chol have non-conforming size")
  if (identical(unique(mean), 0)) return(obs)
  if (length(mean) == J)
    return(obs - c(mean))
  if (!is.matrix(mean))
    stop("obs and mean have non-conforming size")
  if (nrow(mean) != nr)
    stop("obs and mean have non-conforming size")
  if (ncol(mean) != nc)
    stop("obs and mean have non-conforming size")
  return(obs - mean)
}
\diamond
```

Fragment referenced in [2](#).

With  $\boldsymbol{\Sigma}_i = \mathbf{C}_i \mathbf{C}_i^\top$  the log-likelihood function for  $\mathbf{Y}_i = \mathbf{y}_i$  is

$$\ell_i(\boldsymbol{\mu}_i, \mathbf{C}_i) = -\frac{k}{2} \log(2\pi) - \frac{1}{2} \log |\boldsymbol{\Sigma}_i| - \frac{1}{2} (\mathbf{y}_i - \boldsymbol{\mu}_i)^\top \boldsymbol{\Sigma}_i^{-1} (\mathbf{y}_i - \boldsymbol{\mu}_i)$$

Because  $\log |\Sigma_i| = \log |\mathbf{C}_i \mathbf{C}_i^\top| = 2 \log |\mathbf{C}_i| = 2 \sum_{j=1}^J \log \text{diag}(\mathbf{C}_i)_j$  we get the simpler expression

$$\ell_i(\boldsymbol{\mu}_i, \mathbf{C}_i) = -\frac{k}{2} \log(2\pi) - \sum_{j=1}^J \log \text{diag}(\mathbf{C}_i)_j - \frac{1}{2} (\mathbf{y}_i - \boldsymbol{\mu}_i)^\top \mathbf{C}_i^{-\top} \mathbf{C}_i^{-1} (\mathbf{y}_i - \boldsymbol{\mu}_i). \quad (2.1)$$

We need to compute `colSums(dnorm(z, log = TRUE))` quite often. This turns out to be time-consuming and memory intensive, so we provide a small internal helper function focusing on the necessary computations.

`< colSumsdnorm 58a > ≡`

```
SEXP R_ltMatrices_colSumsdnorm (SEXP z, SEXP N, SEXP J) {
  /* number of columns */
  int iN = INTEGER(N)[0];
  /* number of rows */
  int iJ = INTEGER(J)[0];
  SEXP ans;
  double *dans, J12pi, *dz;

  J12pi = iJ * log(2 * M_PI);
  PROTECT(ans = allocVector(REALSXP, iN));
  dans = REAL(ans);
  dz = REAL(z);

  for (int i = 0; i < iN; i++) {
    dans[i] = 0.0;
    for (int j = 0; j < iJ; j++)
      dans[i] += pow(dz[j], 2);
    dans[i] = - 0.5 * (J12pi + dans[i]);
    dz += iJ;
  }

  UNPROTECT(1);
  return(ans);
}
◇
```

Fragment referenced in [3](#).

`< colSumsdnorm ltMatrices 58b > ≡`

```
.colSumsdnorm <- function(z) {
  stopifnot(is.numeric(z))
  if (!is.matrix(z))
    z <- matrix(z, nrow = 1, ncol = length(z))
  ret <- .Call(mvtnorm_R_ltMatrices_colSumsdnorm, z, ncol(z), nrow(z))
  names(ret) <- colnames(z)
  return(ret)
}
◇
```

Fragment referenced in [2](#).

The main part is now

$\langle \text{ldmvnorm chol 59a} \rangle \equiv$

```

if (missing(chol))
  stop("either chol or invchol must be given")
## chol is given
if (!is.ltMatrices(chol))      ### NOTE: replace with is.chol
  stop("chol is not an object of class ltMatrices")
N <- dim(chol)[1L]
N <- ifelse(N == 1, p, N)
J <- dim(chol)[2L]
obs <- .check_obs(obs = obs, mean = mean, J = J, N = N)
z <- solve(chol, obs)
logretval <- .colSumsdnorm(z)
if (attr(chol, "diag"))
  logretval <- logretval - logdet(chol)
◇

```

Fragment referenced in [57a](#).

where we can use the efficient implementations of `solve` and `logdet`.

If  $\mathbf{L}_i = \mathbf{C}_i^{-1}$  is given, we obtain

$$\ell_i(\boldsymbol{\mu}_i, \mathbf{L}_i) = -\frac{k}{2} \log(2\pi) + \sum_{j=1}^J \log \text{diag}(\mathbf{L}_i)_j - \frac{1}{2} (\mathbf{y}_i - \boldsymbol{\mu}_i)^\top \mathbf{L}_i^\top \mathbf{L}_i (\mathbf{y}_i - \boldsymbol{\mu}_i).$$

$\langle \text{ldmvnorm invchol 59b} \rangle \equiv$

```

## invchol is given
if (!is.ltMatrices(invchol))    ### NOTE: replace with is.invchol
  stop("invchol is not an object of class ltMatrices")
N <- dim(invchol)[1L]
N <- ifelse(N == 1, p, N)
J <- dim(invchol)[2L]
obs <- .check_obs(obs = obs, mean = mean, J = J, N = N)
## NOTE: obs is (J x N)
## dnorm takes rather long
z <- Mult(invchol, obs)
logretval <- .colSumsdnorm(z)
## note that the second summand gets recycled the correct number
## of times in case dim(invchol)[1L] == 1 but ncol(obs) > 1
if (attr(invchol, "diag"))
  logretval <- logretval + logdet(invchol)
◇

```

Fragment referenced in [57a](#).

The score function with respect to `obs` is

$$\frac{\partial \ell_i(\boldsymbol{\mu}_i, \mathbf{L}_i)}{\partial \mathbf{y}_i} = -\mathbf{L}_i^\top \mathbf{L}_i \mathbf{y}_i$$

and with respect to `invchol` we have

$$\frac{\partial \ell_i(\boldsymbol{\mu}_i, \mathbf{L}_i)}{\partial \mathbf{L}_i} = -2\mathbf{L}_i \mathbf{y}_i \mathbf{y}_i^\top + \text{diag}(\mathbf{L}_i)^{-1}.$$

The score function with respect to `chol` post-processes the above score using the vec trick (Section 2.11). For the log-likelihood (2.1), the score with respect to  $\mathbf{C}_i$  is the sum of the score functions of the two terms. We start with the simpler first term

$$\frac{\partial - \sum_{j=1}^J \log \text{diag}(\mathbf{C}_i)_j}{\partial \mathbf{C}_i} = -\text{diag}(\mathbf{C}_i)^{-1}$$

The second term gives (we omit the mean for the sake of simplicity)

$$\begin{aligned} \frac{\partial - \mathbf{y}_i^\top \mathbf{C}_i^{-\top} \mathbf{C}_i^{-1} \mathbf{y}_i}{\partial \mathbf{C}_i} &= - \left. \frac{\partial \mathbf{y}_i^\top \mathbf{A}^\top \mathbf{A} \mathbf{y}_i}{\partial \mathbf{A}} \right|_{\mathbf{A}=\mathbf{C}_i^{-1}} \left. \frac{\partial \mathbf{A}^{-1}}{\partial \mathbf{A}} \right|_{\mathbf{A}=\mathbf{C}_i} \\ &= -2\text{vec}(\mathbf{C}_i^{-1} \mathbf{y}_i \mathbf{y}_i^\top)^\top (-1) (\mathbf{C}_i^{-\top} \otimes \mathbf{C}_i^{-1}) \\ &= 2\text{vec}(\mathbf{C}_i^{-\top} \mathbf{C}_i^{-1} \mathbf{y}_i \mathbf{y}_i^\top \mathbf{C}_i^{-\top})^\top \end{aligned}$$

In `sldmvnorm`, we compute the score with respect to  $\mathbf{L}_i$  and use the above relationship to compute the score with respect to  $\mathbf{C}_i$ .

$\langle \text{sldmvnorm 61} \rangle \equiv$

```
sldmvnorm <- function(obs, mean = 0, chol, invchol, logLik = TRUE) {

  stopifnot(xor(missing(chol), missing(invchol)))
  if (!is.matrix(obs)) obs <- matrix(obs, ncol = 1L)

  if (!missing(invchol)) {

    N <- dim(invchol)[1L]
    N <- ifelse(N == 1, ncol(obs), N)
    J <- dim(invchol)[2L]
    obs <- .check_obs(obs = obs, mean = mean, J = J, N = N)

    Mix <- Mult(invchol, obs)
    sobs <- - Mult(invchol, Mix, transpose = TRUE)

    Y <- matrix(obs, byrow = TRUE, nrow = J, ncol = N * J)
    ret <- - matrix(Mix[, rep(1:N, each = J)] * Y, ncol = N)

    M <- matrix(1:(J^2), nrow = J, byrow = FALSE)
    ret <- ret[M[lower.tri(M, diag = attr(invchol, "diag"))],,drop = FALSE]
    if (!is.null(dimnames(invchol)[[1L]]))
      colnames(ret) <- dimnames(invchol)[[1]]
    ret <- ltMatrices(ret,
                      diag = attr(invchol, "diag"), byrow = FALSE,
                      names = dimnames(invchol)[[2L]])
    ret <- ltMatrices(ret, diag = attr(invchol, "diag"),
                      byrow = attr(invchol, "byrow"))
    if (attr(invchol, "diag")) {
      ### recycle properly
      diagonals(ret) <- diagonals(ret) + c(1 / diagonals(invchol))
    } else {
      diagonals(ret) <- 0
    }
    ret <- list(obs = sobs, invchol = ret)
    if (logLik)
      ret$logLik <- ldmvnorm(obs = obs, mean = mean,
                           invchol = invchol, logLik = FALSE)
    return(ret)
  }

  invchol <- solve(chol)
  ret <- sldmvnorm(obs = obs, mean = mean, invchol = invchol)
  ### this means: ret$chol <- - vectrick(invchol, ret$invchol, invchol)
  ret$chol <- as.chol(- vectrick(invchol, ret$invchol))
  ret$invchol <- NULL
  return(ret)
}
◇
```

Fragment referenced in [64](#).

## 2.15 Application Example

Let's say we have  $\mathbf{Y}_i \sim \mathbb{N}_J(\mathbf{0}_J, \mathbf{C}_i \mathbf{C}_i^\top)$  for  $i = 1, \dots, N$  and we know the Cholesky factors  $\mathbf{L}_i = \mathbf{C}_i^{-1}$  of the  $N$  precision matrices  $\Sigma_i^{-1} = \mathbf{L}_i \mathbf{L}_i^\top$ . We generate  $\mathbf{Y}_i = \mathbf{L}_i^{-1} \mathbf{Z}_i$  from  $\mathbf{Z}_i \sim \mathbb{N}_J(\mathbf{0}_J, \mathbf{I}_J)$ .

Evaluating the corresponding log-likelihood is now straightforward and fast, compared to repeated calls to `dmvnorm`

```
> N <- 1000L
> J <- 50L
> lt <- ltMatrices(matrix(runif(N * J * (J + 1) / 2) + 1, ncol = N),
+                   diag = TRUE, byrow = FALSE)
> Z <- matrix(rnorm(N * J), ncol = N)
> Y <- solve(lt, Z)
> l11 <- sum(dnorm(Mult(lt, Y), log = TRUE)) + sum(log(diagonals(lt)))
> S <- as.array(Tcrossprod(solve(lt)))
> l12 <- sum(sapply(1:N, function(i)
+           dmvnorm(x = Y[,i], sigma = S[,i], log = TRUE)))
> chk(l11, l12)
```

The `ldmvnorm` function now also has `chol` and `invchol` arguments such that we can use

```
> l13 <- ldmvnorm(obs = Y, invchol = lt)
> chk(l11, l13)
```

Note that argument `obs` in `ldmvnorm` is an  $J \times N$  matrix whereas the traditional interface in `dmvnorm` expects an  $N \times J$  matrix `x`. The reason is that `Mult` or `solve` work with  $J \times N$  matrices and we want to avoid matrix transposes.

Sometimes it is preferable to split the joint distribution into a marginal distribution of some elements and the conditional distribution given these elements. The joint density is, of course, the product of the marginal and conditional densities and we can check if this works for our example by

```
> ## marginal of and conditional on these
> (j <- 1:5 * 10)

[1] 10 20 30 40 50

> md <- marg_mvnorm(invchol = lt, which = j)
> cd <- cond_mvnorm(invchol = lt, which = j, given = Y[j,])
> l13 <- sum(dnorm(Mult(md$invchol, Y[j,]), log = TRUE)) +
+         sum(log(diagonals(md$invchol))) +
+         sum(dnorm(Mult(cd$invchol, Y[-j,] - cd$mean), log = TRUE)) +
+         sum(log(diagonals(cd$invchol)))
> chk(l11, l13)
```

## Chapter 3

# Multivariate Normal Log-likelihoods

We now discuss code for evaluating the log-likelihood

$$\sum_{i=1}^N \log(p_i(\mathbf{C}_i \mid \mathbf{a}_i, \mathbf{b}_i))$$

This is relatively simple to achieve using the existing `pmvnorm` function, so a prototype might look like

*⟨ lpmvnormR 63 ⟩*  $\equiv$

```
lpmvnormR <- function(lower, upper, mean = 0, center = NULL, chol, logLik = TRUE, ...) {  
  ⟨ input checks 66 ⟩  
  
  sigma <- Tcrossprod(chol)  
  S <- as.array(sigma)  
  idx <- 1  
  
  ret <- error <- numeric(N)  
  for (i in 1:N) {  
    if (dim(sigma)[[1L]] > 1) idx <- i  
    tmp <- pmvnorm(lower = lower[,i], upper = upper[,i], sigma = S[,idx], ...)  
    ret[i] <- tmp  
    error[i] <- attr(tmp, "error")  
  }  
  attr(ret, "error") <- error  
  
  if (logLik)  
    return(sum(log(pmax(ret, .Machine$double.eps))))  
  
  ret  
}
```

◇

Fragment never referenced.

However, the underlying FORTRAN code first computes the Cholesky factor based on the covariance matrix, which is clearly a waste of time. Repeated calls to FORTRAN also cost some time. The code (based on and evaluated in [Genz and Bretz, 2002](#)) implements a specific form of quasi-Monte-Carlo integration without allowing the user to change the scheme (or to fall-back to simple Monte-Carlo). We therefore implement our own simplified version, with the aim to speed-things up such that maximum-likelihood estimation becomes a bit faster.

Let's look at an example first. This code estimates  $p_1, \dots, p_{10}$  for a 5-dimensional normal

```
> J <- 5L
> N <- 10L
> x <- matrix(runif(N * J * (J + 1) / 2), ncol = N)
> lx <- ltMatrices(x, byrow = TRUE, diag = TRUE)
> a <- matrix(runif(N * J), nrow = J) - 2
> a[sample(J * N)[1:2]] <- -Inf
> b <- a + 2 + matrix(runif(N * J), nrow = J)
> b[sample(J * N)[1:2]] <- Inf
> (phat <- c(lpmvnormR(a, b, chol = lx, logLik = FALSE)))

[1] 0.2369329 0.2337179 0.2842052 0.3915213 0.4662496 0.0000000 0.5900784
[8] 0.4618524 0.4872819 0.0000000
```

We want to achieve the same result a bit more general and a bit faster, by making the code more modular and, most importantly, by providing score functions for all arguments  $\mathbf{a}_i$ ,  $\mathbf{b}_i$ , and  $\mathbf{C}_i$ .

### 3.1 Algorithm

"lpmvnorm.R" 64≡

```
< R Header 131 >
< lpmvnorm 75 >
< slpmvnorm 87 >
< ldmvnorm 57a >
< sldmvnorm 61 >
< ldpmvnorm 100 >
< sldpmvnorm 102 >
< deperma 107 >
< standardize 109 >
< destandardize 111 >
◇
```

"lpmvnorm.c" 65≡

```
< C Header 132 >
#ifndef USE_FC_LEN_T
# define USE_FC_LEN_T
#endif
#include <Rconfig.h>
#include <R_ext/BLAS.h> /* for dtrmm */
#ifndef FCONE
# define FCONE
#endif
#include <R.h>
#include <Rmath.h>
#include <Rinternals.h>
#include <Rdefines.h>
< pnorm fast 70b >
< pnorm slow 70c >
< R lpmvnorm 73 >
< R slpmvnorm 84 >
◇
```

We implement the algorithm described by [Genz \(1992\)](#). The key point here is that the original  $J$ -dimensional problem (1.1) is transformed into an integral over  $[0, 1]^{J-1}$ .

For each  $i = 1, \dots, N$ , do

1. Input  $\mathbf{C}_i$  (chol),  $\mathbf{a}_i$  (lower),  $\mathbf{b}_i$  (upper), and control parameters  $\alpha$ ,  $\epsilon$ , and  $M_{\max}$  ( $\mathbf{M}$ ).

$\langle \text{input checks 66} \rangle \equiv$

```

if (!is.matrix(lower)) lower <- matrix(lower, ncol = 1)
if (!is.matrix(upper)) upper <- matrix(upper, ncol = 1)
stopifnot(isTRUE(all.equal(dim(lower), dim(upper))))

stopifnot(is.ltMatrices(chol))          ### NOTE: replace with is.chol
byrow_orig <- attr(chol, "byrow")
chol <- ltMatrices(chol, byrow = TRUE)
d <- dim(chol)
### allow single matrix C
N <- ifelse(d[1L] == 1, ncol(lower), d[1L])
J <- d[2L]

stopifnot(nrow(lower) == J && ncol(lower) == N)
stopifnot(nrow(upper) == J && ncol(upper) == N)
if (is.matrix(mean)) {
  if (ncol(mean) == 1L)
    mean <- mean[,rep(1, N),drop = FALSE]
  stopifnot(nrow(mean) == J && ncol(mean) == N)
}

lower <- lower - mean
upper <- upper - mean

if (!is.null(center)) {
  if (!is.matrix(center)) center <- matrix(center, ncol = 1)
  stopifnot(nrow(center) == J && ncol(center) == N)
}

```

Fragment referenced in [63](#), [75](#), [87](#).

2. Standardise integration limits  $a_j^{(i)}/c_{jj}^{(i)}$ ,  $b_j^{(i)}/c_{jj}^{(i)}$ , and rows  $c_{jj}^{(i)}/c_{jj}^{(i)}$  for  $1 \leq j < j < J$ .

$\langle \text{standardise 67a} \rangle \equiv$

```

if (attr(chol, "diag")) {
  ### diagonals returns J x N and lower/upper are J x N, so
  ### elementwise standardisation is simple
  dchol <- diagonals(chol)
  ### zero diagonals not allowed
  stopifnot(all(abs(dchol) > (.Machine$double.eps)))
  ac <- lower / c(dchol)
  bc <- upper / c(dchol)
  C <- Dchol(chol, D = 1 / dchol)
  if (J > 1) { ### else: univariate problem; C is no longer used
    uC <- Lower_tri(C)
  } else {
    uC <- unclass(C)
  }
} else {
  ac <- lower
  bc <- upper
  uC <- Lower_tri(chol)
}

```

Fragment referenced in [75](#), [87](#).

3. Initialise  $\text{intsum} = \text{varsum} = 0$ ,  $M = 0$ ,  $d_1 = \Phi\left(a_1^{(i)}\right)$ ,  $e_1 = \Phi\left(b_1^{(i)}\right)$  and  $f_1 = e_1 - d_1$ .

$\langle \text{initialisation 67b} \rangle \equiv$

```

x0 = 0.0;
if (LENGTH(center))
  x0 = -dcenter[0];
d0 = pnorm_ptr(da[0], x0);
e0 = pnorm_ptr(db[0], x0);
emd0 = e0 - d0;
f0 = emd0;
intsum = (iJ > 1 ? 0.0 : f0);

```

Fragment referenced in [73](#), [84](#).

4. Repeat

$\langle \text{init logLik loop 67c} \rangle \equiv$

```

d = d0;
f = f0;
emd = emd0;
start = 0;

```

Fragment referenced in [73](#), [79b](#).

- (a) Generate uniform  $w_1, \dots, w_{J-1} \in [0, 1]$ .

(b) For  $j = 2, \dots, J$  set

$$y_{j-1} = \Phi^{-1}(d_{j-1} + w_{j-1}(e_{j-1} - d_{j-1}))$$

We either generate  $w_{j-1}$  on the fly or use pre-computed weights ( $\mathbf{w}$ ).

$\langle \text{compute } y \text{ 68a} \rangle \equiv$

```
Wtmp = (W == R_NilValue ? unif_rand() : dW[j - 1]);
tmp = d + Wtmp * emd;
if (tmp < dtol) {
  y[j - 1] = q0;
} else {
  if (tmp > mdtol)
    y[j - 1] = -q0;
  else
    y[j - 1] = qnorm(tmp, 0.0, 1.0, 1L, 0L);
}
◇
```

Fragment referenced in [69b](#), [83a](#).

$$x_{j-1} = \sum_{j=1}^{j-1} c_{jj}^{(i)} y_j$$

$\langle \text{compute } x \text{ 68b} \rangle \equiv$

```
x = 0.0;
if (LENGTH(center)) {
  for (k = 0; k < j; k++)
    x += dC[start + k] * (y[k] - dcenter[k]);
  x -= dcenter[j];
} else {
  for (k = 0; k < j; k++)
    x += dC[start + k] * y[k];
}
◇
```

Fragment referenced in [69b](#), [83a](#).

$$\begin{aligned} d_j &= \Phi\left(a_j^{(i)} - x_{j-1}\right) \\ e_j &= \Phi\left(b_j^{(i)} - x_{j-1}\right) \end{aligned}$$

$\langle \text{update } d, e \text{ 68c} \rangle \equiv$

```
d = pnorm_ptr(da[j], x);
e = pnorm_ptr(db[j], x);
emd = e - d;
◇
```

Fragment referenced in [69b](#), [83a](#).

$$f_j = (e_j - d_j)f_{j-1}.$$

$\langle \text{update } f \text{ 69a} \rangle \equiv$

```
start += j;
f *= emd;
◇
```

Fragment referenced in [69b](#), [83a](#).

We put everything together in a loop starting with the second dimension

$\langle \text{inner logLik loop 69b} \rangle \equiv$

```
for (j = 1; j < iJ; j++) {
     $\langle \text{compute } y \text{ 68a} \rangle$ 
     $\langle \text{compute } x \text{ 68b} \rangle$ 
     $\langle \text{update } d, e \text{ 68c} \rangle$ 
     $\langle \text{update } f \text{ 69a} \rangle$ 
}
◇
```

Fragment referenced in [73](#).

- (c) Set  $\text{intsum} = \text{intsum} + f_J$ ,  $\text{varsum} = \text{varsum} + f_J^2$ ,  $M = M + 1$ , and  $\text{error} = \sqrt{(\text{varsum}/M - (\text{intsum}/M)^2)/M}$ .

$\langle \text{increment 69c} \rangle \equiv$

```
intsum += f;
◇
```

Fragment referenced in [73](#).

We refrain from early stopping and error estimation.

Until  $\text{error} < \epsilon$  or  $M = M_{\max}$

5. Output  $\hat{p}_i = \text{intsum}/M$ .

We return  $\log \hat{p}_i$  for each  $i$ , or we immediately sum-up over  $i$ .

$\langle \text{output 69d} \rangle \equiv$

```
dans[0] += (intsum < dtol ? 10 : log(intsum)) - 1M;
if (!RlogLik)
    dans += 1L;
◇
```

Fragment referenced in [73](#).

and move on to the next observation (note that  $\mathbf{p}$  might be 0 in case  $\mathbf{C}_i \equiv \mathbf{C}$ ).

$\langle \text{move on 70a} \rangle \equiv$

```

da += iJ;
db += iJ;
dC += p;
if (LENGTH(center)) dcenter += iJ;
◇

```

Fragment referenced in 73, 84.

It turned out that calls to `pnorm` are expensive, so a slightly faster alternative (suggested by [Matić et al., 2018](#)) might provide an alternative which can be requested from using (`fast = TRUE` in the calls to `lpmvnorm` and `slpmvnorm`):

$\langle \text{pnorm fast 70b} \rangle \equiv$

```

/* see https://doi.org/10.2139/ssrn.2842681 */
const double g2 = -0.0150234471495426236132;
const double g4 = 0.000666098511701018747289;
const double g6 = 5.07937324518981103694e-06;
const double g8 = -2.92345273673194627762e-06;
const double g10 = 1.34797733516989204361e-07;
const double m2dpi = -2.0 / M_PI; //3.141592653589793115998;

double C_pnorm_fast (double x, double m) {

    double tmp, ret;
    double x2, x4, x6, x8, x10;

    if (R_FINITE(x)) {
        x = x - m;
        x2 = x * x;
        x4 = x2 * x2;
        x6 = x4 * x2;
        x8 = x6 * x2;
        x10 = x8 * x2;
        tmp = 1 + g2 * x2 + g4 * x4 + g6 * x6 + g8 * x8 + g10 * x10;
        tmp = m2dpi * x2 * tmp;
        ret = .5 + ((x > 0) - (x < 0)) * sqrt(1 - exp(tmp)) / 2.0;
    } else {
        ret = (x > 0 ? 1.0 : 0.0);
    }
    return(ret);
}
◇

```

Fragment referenced in 65.

$\langle \text{pnorm slow 70c} \rangle \equiv$

```

double C_pnorm_slow (double x, double m) {
    return(pnorm(x, m, 1.0, 1L, 0L));
}
◇

```

Fragment referenced in 65.

The `fast` argument can be used to switch on the faster but less accurate version of `pnorm`

$\langle \text{pnorm 71a} \rangle \equiv$

```
Rboolean Rfast = asLogical(fast);
double (*pnorm_ptr)(double, double) = C_pnorm_slow;
if (Rfast)
    pnorm_ptr = C_pnorm_fast;
```

◇

Fragment referenced in [73](#), [84](#).

We allow a new set of weights for each observation or one set for all observations. In the former case, the number of columns is  $M \times N$  and in the latter just  $M$ .

$\langle W \text{ length 71b} \rangle \equiv$

```
int pW = 0;
if (W != R_NilValue) {
    if (LENGTH(W) == (iJ - 1) * iM) {
        pW = 0;
    } else {
        if (LENGTH(W) != (iJ - 1) * iN * iM)
            error("Length of W incorrect");
        pW = 1;
    }
    dW = REAL(W);
}
```

◇

Fragment referenced in [73](#), [84](#).

$\langle \text{dimensions 71c} \rangle \equiv$

```
int iM = INTEGER(M)[0];
int iN = INTEGER(N)[0];
int iJ = INTEGER(J)[0];

da = REAL(a);
db = REAL(b);
dC = REAL(C);
dW = REAL(C); // make -Wmaybe-uninitialized happy

if (LENGTH(C) == iJ * (iJ - 1) / 2)
    p = 0;
else
    p = LENGTH(C) / iN;
```

◇

Fragment referenced in [73](#), [84](#).

$\langle \text{setup return object 72a} \rangle \equiv$

```
len = (RlogLik ? 1 : iN);
PROTECT(ans = allocVector(REALSXP, len));
dans = REAL(ans);
for (int i = 0; i < len; i++)
    dans[i] = 0.0;
◇
```

Fragment referenced in [73](#).

The case  $J = 1$  does not loop over  $M$

$\langle \text{univariate problem 72b} \rangle \equiv$

```
if (iJ == 1) {
    iM = 0;
    lM = 0.0;
} else {
    lM = log((double) iM);
}
◇
```

Fragment referenced in [73](#).

$\langle \text{init center 72c} \rangle \equiv$

```
dcenter = REAL(center);
if (LENGTH(center)) {
    if (LENGTH(center) != iN * iJ)
        error("incorrect dimensions of center");
}
◇
```

Fragment referenced in [73](#), [84](#).

We put the code together in a dedicated C function

$\langle R \text{ slpmvnorm variables 72d} \rangle \equiv$

```
SEXP ans;
double *da, *db, *dC, *dW, *dans, dtol = REAL(tol)[0];
double *dcenter;
double mdtol = 1.0 - dtol;
double d0, e0, emd0, f0, q0;
◇
```

Fragment referenced in [73](#), [84](#).

$\langle R \text{ lpmvnorm } 73 \rangle \equiv$

```
SEXP R_lpmvnorm(SEXP a, SEXP b, SEXP C, SEXP center, SEXP N, SEXP J,
                SEXP W, SEXP M, SEXP tol, SEXP logLik, SEXP fast) {
```

```
   $\langle R \text{ slpmvnorm variables } 72d \rangle$ 
  double l0, lM, x0, intsum;
  int p, len;
```

```
  Rboolean RlogLik = asLogical(logLik);
```

```
   $\langle pnorm \text{ } 71a \rangle$ 
   $\langle dimensions \text{ } 71c \rangle$ 
   $\langle W \text{ length } 71b \rangle$ 
   $\langle init \text{ center } 72c \rangle$ 
```

```
  int start, j, k;
  double tmp, Wtmp, e, d, f, emd, x, y[(iJ > 1 ? iJ - 1 : 1)];
```

```
   $\langle setup \text{ return object } 72a \rangle$ 
```

```
  q0 = qnorm(dtol, 0.0, 1.0, 1L, 0L);
  l0 = log(dtol);
```

```
   $\langle univariate \text{ problem } 72b \rangle$ 
```

```
  if (W == R_NilValue)
    GetRNGstate();
```

```
  for (int i = 0; i < iN; i++) {
```

```
    x0 = 0;
     $\langle initialisation \text{ } 67b \rangle$ 
```

```
    if (W != R_NilValue && pW == 0)
      dW = REAL(W);
```

```
    for (int m = 0; m < iM; m++) {
```

```
       $\langle init \text{ logLik loop } 67c \rangle$ 
       $\langle inner \text{ logLik loop } 69b \rangle$ 
       $\langle increment \text{ } 69c \rangle$ 
```

```
      if (W != R_NilValue)
        dW += iJ - 1;
```

```
    }
```

```
     $\langle output \text{ } 69d \rangle$ 
     $\langle move \text{ on } 70a \rangle$ 
```

```
  }
```

```
  if (W == R_NilValue)
    PutRNGstate();
```

```
  UNPROTECT(1);
  return(ans);
```

```
  }
  ◇
```

Fragment referenced in [65](#).

The R user interface consists of some checks and a call to C. Note that we need to specify both  $\mathbf{w}$  and  $\mathbf{M}$  in case we want a new set of weights for each observation.

*⟨ init random seed, reset on exit 74a ⟩*  $\equiv$

```
### from stats::simulate.lm
if (!exists(".Random.seed", envir = .GlobalEnv, inherits = FALSE))
  runif(1)
if (is.null(seed))
  RNGstate <- get(".Random.seed", envir = .GlobalEnv)
else {
  R.seed <- get(".Random.seed", envir = .GlobalEnv)
  set.seed(seed)
  RNGstate <- structure(seed, kind = as.list(RNGkind()))
  on.exit(assign(".Random.seed", R.seed, envir = .GlobalEnv))
}
◇
```

Fragment referenced in [75](#), [87](#).

*⟨ check and / or set integration weights 74b ⟩*  $\equiv$

```
if (!is.null(w) && J > 1) {
  stopifnot(is.matrix(w))
  stopifnot(nrow(w) == J - 1)
  if (is.null(M))
    M <- ncol(w)
  stopifnot(ncol(w) %in% c(M, M * N))
  if (!is.double(w)) storage.mode(w) <- "double"
} else {
  if (J > 1) {
    if (is.null(M)) stop("either w or M must be specified")
  } else {
    M <- 1L
  }
}
}
◇
```

Fragment referenced in [75](#), [87](#).

Sometimes we want to evaluate the log-likelihood based on  $\mathbf{L} = \mathbf{C}^{-1}$ , the inverse Cholesky factor of the covariance matrix. In this case, we explicitly invert  $\mathbf{L}$  to give  $\mathbf{C}$  (both matrices are lower triangular, so this is fast).

*⟨ Cholesky of precision 74c ⟩*  $\equiv$

```
stopifnot(xor(missing(chol), missing(invchol)))
if (missing(chol)) chol <- solve(invchol)
◇
```

Fragment referenced in [75](#), [87](#).

$\langle \text{lpmvnorm } 75 \rangle \equiv$

```
lpmvnorm <- function(lower, upper, mean = 0, center = NULL, chol, invchol,
                     logLik = TRUE, M = NULL, w = NULL, seed = NULL,
                     tol = .Machine$double.eps, fast = FALSE) {

  < init random seed, reset on exit 74a >
  < Cholesky of precision 74c >
  < input checks 66 >
  < standardise 67a >
  < check and / or set integration weights 74b >

  ret <- .Call(mvtnorm_R_lpmvnorm, ac, bc, uC, as.double(center),
              as.integer(N), as.integer(J), w, as.integer(M), as.double(tol),
              as.logical(logLik), as.logical(fast));
  return(ret)
}
◇
```

Fragment referenced in 64.

Coming back to our simple example, we get (with 25000 simple Monte-Carlo iterations)

```
> phat

[1] 0.2369329 0.2337179 0.2842052 0.3915213 0.4662496 0.0000000 0.5900784
[8] 0.4618524 0.4872819 0.0000000

> exp(lpmvnorm(a, b, chol = lx, M = 25000, logLik = FALSE, fast = TRUE))

[1] 2.366926e-01 2.341369e-01 2.834803e-01 3.938926e-01 4.658150e-01
[6] 8.881784e-21 5.911462e-01 4.597514e-01 4.879485e-01 8.881784e-21

> exp(lpmvnorm(a, b, chol = lx, M = 25000, logLik = FALSE, fast = FALSE))

[1] 2.377131e-01 2.372235e-01 2.831741e-01 3.875320e-01 4.659937e-01
[6] 8.881784e-21 5.895400e-01 4.624243e-01 4.871073e-01 8.881784e-21
```

Next we generate some data and compare our implementation to `pmvnorm` using quasi-Monte-Carlo integration. The `pmvnorm` function uses randomised Korobov rules. The experiment here applies generalised Halton sequences. Plain Monte-Carlo (`w = NULL`) will also work but produces more variable results. Results will depend a lot on appropriate choices and it is the user's responsibility to make sure things work as intended. If you are unsure, you should use `pmvnorm` which provides a well-tested configuration.

```
> M <- 10000L
> if (require("qrng", quietly = TRUE)) {
+   ### quasi-Monte-Carlo
+   W <- t(ghalton(M, d = J - 1))
+ } else {
+   ### Monte-Carlo
+   W <- matrix(runif(M * (J - 1)), nrow = J - 1)
+ }
> ### Genz & Bretz, 2002, without early stopping (really?)
> pGB <- lpmvnormR(a, b, chol = lx, logLik = FALSE,
+   algorithm = GenzBretz(maxpts = M, abseps = 0, releps = 0))
> ### Genz 1992 with quasi-Monte-Carlo, fast pnorm
```

```

> pGqf <- exp(lpmvnorm(a, b, chol = lx, w = W, M = M, logLik = FALSE,
+                   fast = TRUE))
> ### Genz 1992, original Monte-Carlo, fast pnorm
> pGf <- exp(lpmvnorm(a, b, chol = lx, w = NULL, M = M, logLik = FALSE,
+                   fast = TRUE))
> ### Genz 1992 with quasi-Monte-Carlo, R::pnorm
> pGqs <- exp(lpmvnorm(a, b, chol = lx, w = W, M = M, logLik = FALSE,
+                   fast = FALSE))
> ### Genz 1992, original Monte-Carlo, R::pnorm
> pGs <- exp(lpmvnorm(a, b, chol = lx, w = NULL, M = M, logLik = FALSE,
+                   fast = FALSE))
> cbind(pGB, pGqf, pGf, pGqs, pGs)

```

	pGB	pGqf	pGf	pGqs	pGs
[1,]	0.2368918	2.369290e-01	2.344954e-01	2.369297e-01	2.360153e-01
[2,]	0.2341507	2.340099e-01	2.319416e-01	2.340103e-01	2.347435e-01
[3,]	0.2841044	2.841303e-01	2.850959e-01	2.841316e-01	2.870079e-01
[4,]	0.3918357	3.921465e-01	3.931626e-01	3.921469e-01	3.904457e-01
[5,]	0.4671062	4.668249e-01	4.678817e-01	4.668242e-01	4.690837e-01
[6,]	0.0000000	2.220446e-20	2.220446e-20	2.220446e-20	2.220446e-20
[7,]	0.5901670	5.902059e-01	5.907621e-01	5.902056e-01	5.929013e-01
[8,]	0.4613023	4.619428e-01	4.611888e-01	4.619434e-01	4.630231e-01
[9,]	0.4872195	4.870317e-01	4.863298e-01	4.870324e-01	4.820740e-01
[10,]	0.0000000	2.220446e-20	2.220446e-20	2.220446e-20	2.220446e-20

The three versions agree nicely. We now check if the code also works for univariate problems

```

> ### test univariate problem
> ### call pmvnorm
> pGB <- lpmvnormR(a[1,,drop = FALSE], b[1,,drop = FALSE], chol = lx[,1],
+                 logLik = FALSE,
+                 algorithm = GenzBretz(maxpts = M, abseps = 0, releps = 0))
> ### call lpmvnorm
> pGq <- exp(lpmvnorm(a[1,,drop = FALSE], b[1,,drop = FALSE], chol = lx[,1],
+                 logLik = FALSE))
> ### ground truth
> ptr <- pnorm(b[1,] / c(unclass(lx[,1]))) - pnorm(a[1,] / c(unclass(lx[,1])))
> cbind(c(ptr), pGB, pGq)

```

	pGB	pGq
[1,]	0.9999758	0.9999758
[2,]	0.6108928	0.6108928
[3,]	0.9076043	0.9076043
[4,]	0.8979932	0.8979932
[5,]	0.9589363	0.9589363
[6,]	0.7863435	0.7863435
[7,]	0.9982537	0.9982537
[8,]	0.8745388	0.8745388
[9,]	0.9386051	0.9386051
[10,]	0.9119778	0.9119778

Because the default `fast = FALSE` was used here, all results are identical.

## 3.2 Score Function

In addition to the log-likelihood, we would also like to have access to the scores with respect to  $\mathbf{C}_i$ . Because every element of  $\mathbf{C}_i$  only enters once, the chain rule rules, so to speak.

We need the derivatives of  $d$ ,  $e$ ,  $y$ , and  $f$  with respect to the  $c$  parameters

$\langle \text{chol scores 77a} \rangle \equiv$

```
double dp_c[Jp], ep_c[Jp], fp_c[Jp], yp_c[(iJ > 1 ? iJ - 1 : 1) * Jp];
◇
```

Fragment referenced in [78a](#).

and the derivates with respect to the mean

$\langle \text{mean scores 77b} \rangle \equiv$

```
double dp_m[Jp], ep_m[Jp], fp_m[Jp], yp_m[(iJ > 1 ? iJ - 1 : 1) * Jp];
◇
```

Fragment referenced in [78a](#).

and the derivates with respect to lower (a)

$\langle \text{lower scores 77c} \rangle \equiv$

```
double dp_l[Jp], ep_l[Jp], fp_l[Jp], yp_l[(iJ > 1 ? iJ - 1 : 1) * Jp];
◇
```

Fragment referenced in [78a](#).

and the derivates with respect to upper (b)

$\langle \text{upper scores 77d} \rangle \equiv$

```
double dp_u[Jp], ep_u[Jp], fp_u[Jp], yp_u[(iJ > 1 ? iJ - 1 : 1) * Jp];
◇
```

Fragment referenced in [78a](#).

and we start allocating the necessary memory. The output object contains the likelihood contributions (first row), the scores with respect to the mean (next  $J$  rows), with respect to the lower integration limits (next  $J$  rows), with respect to the upper integration limits (next  $J$  rows) and finally with respect to the off-diagonal elements of the Cholesky factor (last  $J(J-1)/2$  rows).

$\langle \text{score output object 78a} \rangle \equiv$

```

int Jp = iJ * (iJ + 1) / 2;
 $\langle \text{chol scores 77a} \rangle$ 
 $\langle \text{mean scores 77b} \rangle$ 
 $\langle \text{lower scores 77c} \rangle$ 
 $\langle \text{upper scores 77d} \rangle$ 
double dtmp, etmp, Wtmp, ytmp, xx;

PROTECT(ans = allocMatrix(REALSXP, Jp + 1 + 3 * iJ, iN));
dans = REAL(ans);
for (j = 0; j < LENGTH(ans); j++) dans[j] = 0.0;

```

Fragment referenced in 84.

For each  $i = 1, \dots, N$ , do

1. Input  $\mathbf{C}_i$  (chol),  $\mathbf{a}_i$  (lower),  $\mathbf{b}_i$  (upper), and control parameters  $\alpha$ ,  $\epsilon$ , and  $M_{\max}(\mathbf{M})$ .
2. Standardise integration limits  $a_j^{(i)}/c_{jj}^{(i)}$ ,  $b_j^{(i)}/c_{jj}^{(i)}$ , and rows  $c_{jj}^{(i)}/c_{jj}^{(i)}$  for  $1 \leq j < j < J$ .

Note: We later need derivatives wrt  $c_{jj}^{(i)}$ , so we compute derivatives wrt  $a_j^{(i)}$  and  $b_j^{(i)}$  and post-differentiate later.

3. Initialise  $\text{intsum} = \text{varsum} = 0$ ,  $M = 0$ ,  $d_1 = \Phi(a_1^{(i)})$ ,  $e_1 = \Phi(b_1^{(i)})$  and  $f_1 = e_1 - d_1$ .

We start initialised the score wrt to  $c_{11}^{(i)}$  (the parameter is non-existent here due to standardisation)

$\langle \text{score c11 78b} \rangle \equiv$

```

if (LENGTH(center)) {
  dp_c[0] = (R_FINITE(da[0]) ? dnorm(da[0], x0, 1.0, 0L) * (da[0] - x0 - dcenter[0]) : 0);
  ep_c[0] = (R_FINITE(db[0]) ? dnorm(db[0], x0, 1.0, 0L) * (db[0] - x0 - dcenter[0]) : 0);
} else {
  dp_c[0] = (R_FINITE(da[0]) ? dnorm(da[0], x0, 1.0, 0L) * (da[0] - x0) : 0);
  ep_c[0] = (R_FINITE(db[0]) ? dnorm(db[0], x0, 1.0, 0L) * (db[0] - x0) : 0);
}
fp_c[0] = ep_c[0] - dp_c[0];

```

Fragment referenced in 79b, 84.

$\langle \text{score a, b 79a} \rangle \equiv$

```

dp_m[0] = (R_FINITE(da[0]) ? dnorm(da[0], x0, 1.0, 0L) : 0);
ep_m[0] = (R_FINITE(db[0]) ? dnorm(db[0], x0, 1.0, 0L) : 0);
dp_l[0] = dp_m[0];
ep_u[0] = ep_m[0];
dp_u[0] = 0;
ep_l[0] = 0;
fp_m[0] = ep_m[0] - dp_m[0];
fp_l[0] = -dp_m[0];
fp_u[0] = ep_m[0];

```

Fragment referenced in 79b, 84.

#### 4. Repeat

$\langle \text{init score loop 79b} \rangle \equiv$

$\langle \text{init logLik loop 67c} \rangle$   
 $\langle \text{score c11 78b} \rangle$   
 $\langle \text{score a, b 79a} \rangle$   
 $\diamond$

Fragment referenced in 84.

(a) Generate uniform  $w_1, \dots, w_{J-1} \in [0, 1]$ .

(b) For  $j = 2, \dots, J$  set

$$y_{j-1} = \Phi^{-1}(d_{j-1} + w_{j-1}(e_{j-1} - d_{j-1}))$$

We again either generate  $w_{j-1}$  on the fly or use pre-computed weights (**w**). We first compute the scores with respect to the already existing parameters.

$\langle \text{update yp for chol 79c} \rangle \equiv$

```

ytmp = exp(- dnorm(y[j - 1], 0.0, 1.0, 1L)); // = 1 / dnorm(y[j - 1], 0.0, 1.0, 0L)

for (k = 0; k < Jp; k++) yp_c[k * (iJ - 1) + (j - 1)] = 0.0;

for (idx = 0; idx < (j + 1) * j / 2; idx++) {
  yp_c[idx * (iJ - 1) + (j - 1)] = ytmp;
  yp_c[idx * (iJ - 1) + (j - 1)] *= (dp_c[idx] + Wtmp * (ep_c[idx] - dp_c[idx]));
}

```

$\diamond$

Fragment referenced in 83a.

$\langle \text{update yp for means, lower and upper 80} \rangle \equiv$

```

for (k = 0; k < iJ; k++)
  yp_m[k * (iJ - 1) + (j - 1)] = 0.0;

for (idx = 0; idx < j; idx++) {
  yp_m[idx * (iJ - 1) + (j - 1)] = ytmp;
  yp_m[idx * (iJ - 1) + (j - 1)] *= (dp_m[idx] + Wtmp * (ep_m[idx] - dp_m[idx]));
}

for (k = 0; k < iJ; k++)
  yp_l[k * (iJ - 1) + (j - 1)] = 0.0;

for (idx = 0; idx < j; idx++) {
  yp_l[idx * (iJ - 1) + (j - 1)] = ytmp;
  yp_l[idx * (iJ - 1) + (j - 1)] *= (dp_l[idx] + Wtmp * (dp_u[idx] - dp_l[idx]));
}

for (k = 0; k < iJ; k++)
  yp_u[k * (iJ - 1) + (j - 1)] = 0.0;

for (idx = 0; idx < j; idx++) {
  yp_u[idx * (iJ - 1) + (j - 1)] = ytmp;
  yp_u[idx * (iJ - 1) + (j - 1)] *= (ep_l[idx] + Wtmp * (ep_u[idx] - ep_l[idx]));
}

```

$\diamond$

Fragment referenced in 83a.

$$x_{j-1} = \sum_{j=1}^{j-1} c_{jj}^{(i)} y_j$$

$$\begin{aligned} d_j &= \Phi \left( a_j^{(i)} - x_{j-1} \right) \\ e_j &= \Phi \left( b_j^{(i)} - x_{j-1} \right) \end{aligned}$$

$$f_j = (e_j - d_j) f_{j-1}.$$

The scores with respect to  $c_{jj}^{(i)}, j = 1, \dots, j-1$  are

$\langle \text{score wrt new chol off-diagonals 81a} \rangle \equiv$

```

dtmp = dnorm(da[j], x, 1.0, 0L);
etmp = dnorm(db[j], x, 1.0, 0L);

for (k = 0; k < j; k++) {
  idx = start + j + k;
  if (LENGTH(center)) {
    dp_c[idx] = dtmp * (-1.0) * (y[k] - dcenter[k]);
    ep_c[idx] = etmp * (-1.0) * (y[k] - dcenter[k]);
  } else {
    dp_c[idx] = dtmp * (-1.0) * y[k];
    ep_c[idx] = etmp * (-1.0) * y[k];
  }
  fp_c[idx] = (ep_c[idx] - dp_c[idx]) * f;
}
◇

```

Fragment referenced in [83a](#).

and the score with respect to (the here non-existing)  $c_{jj}^{(i)}$  is

$\langle \text{score wrt new chol diagonal 81b} \rangle \equiv$

```

idx = (j + 1) * (j + 2) / 2 - 1;
if (LENGTH(center)) {
  dp_c[idx] = (R_FINITE(da[j]) ? dtmp * (da[j] - x - dcenter[j]) : 0);
  ep_c[idx] = (R_FINITE(db[j]) ? etmp * (db[j] - x - dcenter[j]) : 0);
} else {
  dp_c[idx] = (R_FINITE(da[j]) ? dtmp * (da[j] - x) : 0);
  ep_c[idx] = (R_FINITE(db[j]) ? etmp * (db[j] - x) : 0);
}
fp_c[idx] = (ep_c[idx] - dp_c[idx]) * f;
◇

```

Fragment referenced in [83a](#).

$\langle \text{new score means, lower and upper 81c} \rangle \equiv$

```

dp_m[j] = (R_FINITE(da[j]) ? dtmp : 0);
ep_m[j] = (R_FINITE(db[j]) ? etmp : 0);
dp_l[j] = dp_m[j];
ep_u[j] = ep_m[j];
dp_u[j] = 0;
ep_l[j] = 0;
fp_l[j] = - dp_m[j] * f;
fp_u[j] = ep_m[j] * f;
fp_m[j] = fp_u[j] + fp_l[j];

```

◇

Fragment referenced in [83a](#).

We next update scores for parameters introduced for smaller  $j$

$\langle \text{update score for chol 82a} \rangle \equiv$

```

for (idx = 0; idx < j * (j + 1) / 2; idx++) {
  xx = 0.0;
  for (k = 0; k < j; k++)
    xx += dC[start + k] * yp_c[idx * (iJ - 1) + k];

  dp_c[idx] = dtmp * (-1.0) * xx;
  ep_c[idx] = etmp * (-1.0) * xx;
  fp_c[idx] = (ep_c[idx] - dp_c[idx]) * f + emd * fp_c[idx];
}

```

◇

Fragment referenced in [83a](#).

$\langle \text{update score means, lower and upper 82b} \rangle \equiv$

```

for (idx = 0; idx < j; idx++) {
  xx = 0.0;
  for (k = 0; k < j; k++)
    xx += dC[start + k] * yp_m[idx * (iJ - 1) + k];

  dp_m[idx] = dtmp * (-1.0) * xx;
  ep_m[idx] = etmp * (-1.0) * xx;
  fp_m[idx] = (ep_m[idx] - dp_m[idx]) * f + emd * fp_m[idx];
}

for (idx = 0; idx < j; idx++) {
  xx = 0.0;
  for (k = 0; k < j; k++)
    xx += dC[start + k] * yp_l[idx * (iJ - 1) + k];

  dp_l[idx] = dtmp * (-1.0) * xx;
  dp_u[idx] = etmp * (-1.0) * xx;
  fp_l[idx] = (dp_u[idx] - dp_l[idx]) * f + emd * fp_l[idx];
}

for (idx = 0; idx < j; idx++) {
  xx = 0.0;
  for (k = 0; k < j; k++)
    xx += dC[start + k] * yp_u[idx * (iJ - 1) + k];

  ep_l[idx] = dtmp * (-1.0) * xx;
  ep_u[idx] = etmp * (-1.0) * xx;
  fp_u[idx] = (ep_u[idx] - ep_l[idx]) * f + emd * fp_u[idx];
}

```

◇

Fragment referenced in [83a](#).

We put everything together in a loop starting with the second dimension

$\langle \text{inner score loop 83a} \rangle \equiv$

```

for (j = 1; j < iJ; j++) {

   $\langle \text{compute } y \text{ 68a} \rangle$ 
   $\langle \text{compute } x \text{ 68b} \rangle$ 
   $\langle \text{update } d, e \text{ 68c} \rangle$ 
   $\langle \text{update } yp \text{ for chol 79c} \rangle$ 
   $\langle \text{update } yp \text{ for means, lower and upper 80} \rangle$ 
   $\langle \text{score wrt new chol off-diagonals 81a} \rangle$ 
   $\langle \text{score wrt new chol diagonal 81b} \rangle$ 
   $\langle \text{new score means, lower and upper 81c} \rangle$ 
   $\langle \text{update score for chol 82a} \rangle$ 
   $\langle \text{update score means, lower and upper 82b} \rangle$ 
   $\langle \text{update } f \text{ 69a} \rangle$ 

}

```

◇

Fragment referenced in [84](#).

(c) Set  $\text{intsum} = \text{intsum} + f_J$ ,  $\text{varsum} = \text{varsum} + f_J^2$ ,  $M = M + 1$ , and  $\text{error} = \sqrt{(\text{varsum}/M - (\text{intsum}/M)^2)/M}$ .

We refrain from early stopping and error estimation.

Until  $\text{error} < \epsilon$  or  $M = M_{\max}$

5. Output  $\hat{p}_i = \text{intsum}/M$ .

We return  $\log \hat{p}_i$  for each  $i$ , or we immediately sum-up over  $i$ .

$\langle \text{score output 83b} \rangle \equiv$

```

dans[0] += f;
for (j = 0; j < Jp; j++)
    dans[j + 1] += fp_c[j];
for (j = 0; j < iJ; j++) {
    idx = Jp + j + 1;
    dans[idx] += fp_m[j];
    dans[idx + iJ] += fp_l[j];
    dans[idx + 2 * iJ] += fp_u[j];
}
◇

```

Fragment referenced in [84](#).

$\langle \text{init dans 83c} \rangle \equiv$

```

if (iM == 0) {
    dans[0] = intsum;
    dans[1] = fp_c[0];
    dans[2] = fp_m[0];
    dans[3] = fp_l[0];
    dans[4] = fp_u[0];
}
◇

```

Fragment referenced in [84](#).

We put everything together in C

$\langle R \text{ slpmvnorm } 84 \rangle \equiv$

```

SEXP R_slpmvnorm(SEXP a, SEXP b, SEXP C, SEXP center, SEXP N, SEXP J, SEXP W,
                 SEXP M, SEXP tol, SEXP fast) {

   $\langle R \text{ slpmvnorm variables } 72d \rangle$ 
  double intsum;
  int p, idx;
   $\langle dimensions \ 71c \rangle$ 
   $\langle pnorm \ 71a \rangle$ 
   $\langle W \text{ length } 71b \rangle$ 
   $\langle init \ center \ 72c \rangle$ 
  int start, j, k;
  double tmp, e, d, f, emd, x, x0, y[(iJ > 1 ? iJ - 1 : 1)];

   $\langle score \ output \ object \ 78a \rangle$ 

  q0 = qnorm(dtol, 0.0, 1.0, 1L, 0L);

  /* univariate problem */
  if (iJ == 1) iM = 0;

  if (W == R_NilValue)
    GetRNGstate();

  for (int i = 0; i < iN; i++) {

     $\langle initialisation \ 67b \rangle$ 
     $\langle score \ c11 \ 78b \rangle$ 
     $\langle score \ a, \ b \ 79a \rangle$ 
     $\langle init \ dans \ 83c \rangle$ 

    if (W != R_NilValue && pW == 0)
      dW = REAL(W);

    for (int m = 0; m < iM; m++) {
       $\langle init \ score \ loop \ 79b \rangle$ 
       $\langle inner \ score \ loop \ 83a \rangle$ 
       $\langle score \ output \ 83b \rangle$ 
      if (W != R_NilValue)
        dW += iJ - 1;
    }

     $\langle move \ on \ 70a \rangle$ 
    dans += Jp + 1 + 3 * iJ;
  }

  if (W == R_NilValue)
    PutRNGstate();

  UNPROTECT(1);
  return(ans);
}

```

◇

Fragment referenced in 65.

The R code is now essentially identical to `lpmvnorm`, however, we need to undo the effect of

standardisation once the scores have been computed

$\langle \text{post differentiate mean score 85a} \rangle \equiv$

```
Jp <- J * (J + 1) / 2;
smean <- - ret[Jp + 1:J, , drop = FALSE]
if (attr(chol, "diag"))
  smean <- smean / c(dchol)
◇
```

Fragment referenced in [87](#).

$\langle \text{post differentiate lower score 85b} \rangle \equiv$

```
slower <- ret[Jp + J + 1:J, , drop = FALSE]
if (attr(chol, "diag"))
  slower <- slower / c(dchol)
◇
```

Fragment referenced in [87](#).

$\langle \text{post differentiate upper score 85c} \rangle \equiv$

```
supper <- ret[Jp + 2 * J + 1:J, , drop = FALSE]
if (attr(chol, "diag"))
  supper <- supper / c(dchol)
◇
```

Fragment referenced in [87](#).

$\langle \text{post differentiate chol score 85d} \rangle \equiv$

```
if (J == 1) {
  idx <- 1L
} else {
  idx <- cumsum(c(1, 2:J))
}
if (attr(chol, "diag")) {
  ret <- ret / c(dchol[rep(1:J, 1:J),]) ### because 1 / dchol already there
  ret[idx,] <- -ret[idx,]
}
◇
```

Fragment referenced in [87](#).

We sometimes parameterise models in terms of  $\mathbf{L} = \mathbf{C}^{-1}$ , the Cholesky factor of the precision matrix. The log-likelihood operates on  $\mathbf{C}$ , so we need to post-differentiate the score function. We have

$$\mathbf{A} = \frac{\partial \mathbf{L}^{-1}}{\partial \mathbf{L}} = -\mathbf{L}^{-\top} \otimes \mathbf{L}^{-1}$$

and computing  $\mathbf{sA}$  for a score vector  $\mathbf{s}$  with respect to  $\mathbf{L}$  can be implemented by the “vec trick” (Section [2.11](#))

$$\mathbf{sA} = \mathbf{L}^{-\top} \mathbf{S} \mathbf{L}^{-\top}$$

where  $\mathbf{s} = \text{vec}(\mathbf{S})$ .

$\langle \text{post differentiate invchol score 86a} \rangle \equiv$

```
if (!missing(invchol)) {  
  ret <- ltMatrices(ret, diag = TRUE, byrow = TRUE,  
                    names = dimnames(chol)[[2L]])  
  ### this means vectrick(chol, ret, chol)  
  ret <- - unclass(vectrick(chol, ret))  
}  
◇
```

Fragment referenced in [87](#).

If the diagonal elements are constants, we set them to zero. The function always returns an object of class `ltMatrices` with explicit diagonal elements (use `Lower_tri(, diag = FALSE)` to extract the lower triangular elements such that the scores match the input)

$\langle \text{post process score 86b} \rangle \equiv$

```
if (!attr(chol, "diag"))  
  ### remove scores for constant diagonal elements  
  ret[idx,] <- 0  
ret <- ltMatrices(ret, diag = TRUE, byrow = TRUE,  
                  names = dimnames(chol)[[2L]])  
◇
```

Fragment referenced in [87](#).

We can now finally put everything together in a single score function.

$\langle \text{slpmvnorm } 87 \rangle \equiv$

```
slpmvnorm <- function(lower, upper, mean = 0, center = NULL,
                      chol, invchol, logLik = TRUE, M = NULL,
                      w = NULL, seed = NULL, tol = .Machine$double.eps,
                      fast = FALSE) {

   $\langle$  init random seed, reset on exit 74a  $\rangle$ 
   $\langle$  Cholesky of precision 74c  $\rangle$ 
   $\langle$  input checks 66  $\rangle$ 
   $\langle$  standardise 67a  $\rangle$ 
   $\langle$  check and / or set integration weights 74b  $\rangle$ 

  ret <- .Call(mvtnorm_R_slpmvnorm, ac, bc, uC, as.double(center), as.integer(N),
              as.integer(J), w, as.integer(M), as.double(tol), as.logical(fast));

  ll <- log(pmax(ret[1L,], tol)) - log(M)
  intsum <- ret[1L,]
  m <- matrix(intsum, nrow = nrow(ret) - 1, ncol = ncol(ret), byrow = TRUE)
  ret <- ret[-1L,,drop = FALSE] / m ### NOTE: division by zero MAY happen,
                                   ### catch outside

   $\langle$  post differentiate mean score 85a  $\rangle$ 
   $\langle$  post differentiate lower score 85b  $\rangle$ 
   $\langle$  post differentiate upper score 85c  $\rangle$ 

  ret <- ret[1:Jp, , drop = FALSE]

   $\langle$  post differentiate chol score 85d  $\rangle$ 
   $\langle$  post differentiate invchol score 86a  $\rangle$ 
   $\langle$  post process score 86b  $\rangle$ 

  ret <- ltMatrices(ret, byrow = byrow_orig)

  rownames(smean) <- rownames(slower) <-
    rownames(supper) <- dimnames(chol)[[2L]]

  if (logLik) {
    ret <- list(logLik = ll,
               mean = smean,
               lower = slower,
               upper = supper,
               chol = ret)
    if (!missing(invchol)) names(ret)[names(ret) == "chol"] <- "invchol"
    return(ret)
  }

  return(ret)
}
◇
```

Fragment referenced in 64.

Let's look at an example, where we use `numDeriv::grad` to check the results (this functionality from package **numDeriv** was absolutely instrumental for this project)

```
> J <- 5L
> N <- 4L
```

```

> S <- crossprod(matrix(runif(J^2), nrow = J))
> prm <- t(chol(S))[lower.tri(S, diag = TRUE)]
> ### define C
> mC <- ltMatrices(matrix(prm, ncol = 1), diag = TRUE)
> a <- matrix(runif(N * J), nrow = J) - 2
> b <- a + 4
> a[2,] <- -Inf
> b[3,] <- Inf
> M <- 10000L
> W <- matrix(runif(M * (J - 1)), ncol = M)
> lli <- c(lpmvnorm(a, b, chol = mC, w = W, M = M, logLik = FALSE))
> fC <- function(prm) {
+   C <- ltMatrices(matrix(prm, ncol = 1), diag = TRUE)
+   lpmvnorm(a, b, chol = C, w = W, M = M)
+ }
> sC <- slpmvnorm(a, b, chol = mC, w = W, M = M)
> chk(lli, sC$logLik)
> if (require("numDeriv", quietly = TRUE))
+   chk(grad(fC, unclass(mC)), rowSums(unclass(sC$chol)),
+       check.attributes = FALSE)

```

We can do the same when **L** (and not **C**) is given

```

> mL <- solve(mC)
> lliL <- c(lpmvnorm(a, b, invchol = mL, w = W, M = M, logLik = FALSE))
> chk(lli, lliL)
> fL <- function(prm) {
+   L <- ltMatrices(matrix(prm, ncol = 1), diag = TRUE)
+   lpmvnorm(a, b, invchol = L, w = W, M = M)
+ }
> sL <- slpmvnorm(a, b, invchol = mL, w = W, M = M)
> chk(lliL, sL$logLik)
> if (require("numDeriv", quietly = TRUE))
+   chk(grad(fL, unclass(mL)), rowSums(unclass(sL$invchol)),
+       check.attributes = FALSE)

```

The score function also works for univariate problems

```

> ptr <- pnorm(b[1,] / c(unclass(mC[,1]))) - pnorm(a[1,] / c(unclass(mC[,1])))
> log(ptr)

[1] -0.01165889 -0.08617272 -0.01240094 -0.03105050

> lpmvnorm(a[1,,drop = FALSE], b[1,,drop = FALSE], chol = mC[,1], logLik = FALSE)

[1] -0.01165889 -0.08617272 -0.01240094 -0.03105050

> lapply(slpvmvnorm(a[1,,drop = FALSE], b[1,,drop = FALSE], chol = mC[,1],
+   logLik = TRUE), unclass)

$logLik
[1] -0.01165889 -0.08617272 -0.01240094 -0.03105050

$mean
      [,1]      [,2]      [,3]      [,4]
1 0.02222249 0.2140162 0.02641782 0.08861162

```

```

$lower
      [,1]      [,2]      [,3]      [,4]
1 -0.03221736 -0.214453 -0.03536199 -0.09096213

$upper
      [,1]      [,2]      [,3]      [,4]
1 0.00999487 0.0004368597 0.008944164 0.002350511

$chol
      [,1]      [,2]      [,3]      [,4]
1.1 -0.104149 -0.2994286 -0.1075726 -0.1787174
attr(,"J")
[1] 1
attr(,"diag")
[1] TRUE
attr(,"byrow")
[1] FALSE
attr(,"rcnames")
[1] "1"

> sd1 <- c(unclass(mC[,1]))
> (dnorm(b[1,] / sd1) * b[1,] - dnorm(a[1,] / sd1) * a[1,]) * (-1) / sd1^2 / ptr

[1] -0.1041490 -0.2994286 -0.1075726 -0.1787174

```

## Chapter 4

# Maximum-likelihood Example

We now discuss how this infrastructure can be used to estimate the Cholesky factor of a multivariate normal in the presence of interval-censored observations.

We first generate a covariance matrix  $\Sigma = \mathbf{C}\mathbf{C}^\top$  and extract the Cholesky factor  $\mathbf{C}$

```
> J <- 4
> R <- diag(J)
> R[1,2] <- R[2,1] <- .25
> R[1,3] <- R[3,1] <- .5
> R[2,4] <- R[4,2] <- .75
> ### ATLAS and M1mac print 0 as something < .Machine$double.eps
> round(Sigma <- diag(sqrt(1:J / 2)) %*% R %*% diag(sqrt(1:J / 2)), 7)
```

```
      [,1]      [,2]      [,3]      [,4]
[1,] 0.5000000 0.1767767 0.4330127 0.000000
[2,] 0.1767767 1.0000000 0.0000000 1.06066
[3,] 0.4330127 0.0000000 1.5000000 0.00000
[4,] 0.0000000 1.0606602 0.0000000 2.00000
```

```
> (C <- t(chol(Sigma)))
```

```
      [,1]      [,2]      [,3]      [,4]
[1,] 0.7071068 0.0000000 0.0000000 0.0000000
[2,] 0.2500000 0.9682458 0.0000000 0.0000000
[3,] 0.6123724 -0.1581139 1.0488088 0.0000000
[4,] 0.0000000 1.0954451 0.1651446 0.8790491
```

We now represent this matrix as `ltMatrices` object

```
> prm <- C[lower.tri(C, diag = TRUE)]
> lt <- ltMatrices(matrix(prm, ncol = 1L),
+                   diag = TRUE,      ### has diagonal elements
+                   byrow = FALSE)    ### prm is column-major
> BYROW <- FALSE    ### later checks
> lt <- ltMatrices(lt,
+                   byrow = BYROW)    ### convert to row-major
> chk(C, as.array(lt)[,,1], check.attributes = FALSE)
> chk(Sigma, as.array(Tcrossprod(lt))[,,1], check.attributes = FALSE)
```

We generate some data from  $\mathbb{N}_J(\mathbf{0}_J, \Sigma)$  by first sampling from  $\mathbf{Z} \sim \mathbb{N}_J(\mathbf{0}_J, \mathbf{I}_J)$  and then computing  $\mathbf{Y} = \mathbf{C}\mathbf{Z} + \boldsymbol{\mu} \sim \mathbb{N}_J(\boldsymbol{\mu}, \mathbf{C}\mathbf{C}^\top)$

```

> N <- 100L
> Z <- matrix(rnorm(N * J), nrow = J)
> Y <- Mult(lt, Z) + (mn <- 1:J)

```

Before we add some interval-censoring to the data, let's estimate the Cholesky factor  $\mathbf{C}$  (here called `lt`) from the raw continuous data. The true mean  $\boldsymbol{\mu}$  and the true covariance matrix  $\Sigma$  can be estimated from the uncensored data via maximum likelihood as

```

> rowMeans(Y)

      1      2      3      4
0.9685377 2.1268796 2.9633561 3.9825669

> (Shat <- var(t(Y)) * (N - 1) / N)

      1      2      3      4
1 0.46655660 0.18104431 0.34222237 0.01609179
2 0.18104431 0.94385339 0.08992252 0.84309528
3 0.34222237 0.08992252 1.36054915 0.08104091
4 0.01609179 0.84309528 0.08104091 1.63301525

```

We first check if we can obtain the same results by numerical optimisation using `dmvnorm` and the scores `sldmvnorm`. The log-likelihood and the score function (for the centered means) in terms of  $\mathbf{C}$  are

```

> Yc <- Y - rowMeans(Y)
> ll <- function(parm) {
+   C <- ltMatrices(parm, diag = TRUE, byrow = BYROW)
+   -ldmvnorm(obs = Yc, chol = C)
+ }
> sc <- function(parm) {
+   C <- ltMatrices(parm, diag = TRUE, byrow = BYROW)
+   -rowSums(unclass(sldmvnorm(obs = Yc, chol = C)$chol))
+ }

```

The diagonal elements of  $\mathbf{C}$  are positive, so we need box constraints

```

> llim <- rep(-Inf, J * (J + 1) / 2)
> llim[which(rownames(unclass(lt)) %in% paste(1:J, 1:J, sep = "."))] <- 1e-4

```

The ML-estimate of  $\mathbf{C}\mathbf{C}^\top$  is now used to obtain an estimate of  $\mathbf{C}$  and we check the score function for some random starting values

```

> if (BYROW) {
+   cML <- chol(Shat)[upper.tri(Shat, diag = TRUE)]
+ } else {
+   cML <- t(chol(Shat))[lower.tri(Shat, diag = TRUE)]
+ }
> ll(cML)

[1] 517.8685

> start <- runif(length(cML))
> if (require("numDeriv", quietly = TRUE))
+   chk(grad(ll, start), sc(start), check.attributes = FALSE)

```

Finally, we hand over to `optim` and compare the results of the analytically and numerically obtained ML estimates

```
> op <- optim(start, fn = ll, gr = sc, method = "L-BFGS-B",
+           lower = llim, control = list(trace = TRUE))
```

```
iter   10 value 518.092239
iter   20 value 517.868548
final  value 517.868548
converged
```

```
> ## ML numerically
> ltMatrices(op$par, diag = TRUE, byrow = BYROW)

, , 1
```

```
      1      2      3      4
1 0.68305690 0.00000000 0.0000000 0.0000000
2 0.26505417 0.93464707 0.0000000 0.0000000
3 0.50102358 -0.04586658 1.0523442 0.0000000
4 0.02356369 0.89534692 0.1048239 0.9054404
```

```
> ll(op$par)
```

```
[1] 517.8685
```

```
> ## ML analytically
> t(chol(Shat))
```

```
      1      2      3      4
1 0.68304949 0.00000000 0.0000000 0.0000000
2 0.26505300 0.93466588 0.0000000 0.0000000
3 0.50102134 -0.04587167 1.052341 0.0000000
4 0.02355875 0.89534773 0.104822 0.9054419
```

```
> ll(cML)
```

```
[1] 517.8685
```

```
> ## true C matrix
```

```
> lt
```

```
, , 1
```

```
      1      2      3      4
1 0.7071068 0.0000000 0.0000000 0.0000000
2 0.2500000 0.9682458 0.0000000 0.0000000
3 0.6123724 -0.1581139 1.0488088 0.0000000
4 0.0000000 1.0954451 0.1651446 0.8790491
```

Under interval-censoring, the mean and **C** are no longer orthogonal and there is no analytic solution to the ML estimation problem. So, we add some interval-censoring represented by **lwr** and **upr** and try to estimate the model parameters via **lpmvnorm** and corresponding scores **slpmvnorm**.

```
> prb <- 1:9 / 10
> sds <- sqrt(diag(Sigma))
> ct <- sapply(1:J, function(j) qnorm(prb, mean = mn[j], sd = sds[j]))
> lwr <- upr <- Y
> for (j in 1:J) {
+   f <- cut(Y[j,], breaks = c(-Inf, ct[,j], Inf))
+   lwr[j,] <- c(-Inf, ct[,j])[f]
+   upr[j,] <- c(ct[,j], Inf)[f]
+ }
```

Let's do some sanity and performance checks first. For different values of  $M$ , we evaluate the log-likelihood using `pmvnorm` (called in `lpmvnormR`) and the simplified implementation (fast and slow). The comparison is a bit unfair, because we do not add the time needed to setup Halton sequences, but we would do this only once and use the stored values for repeated evaluations of a log-likelihood (because the optimiser expects a deterministic function to be optimised)

```
> M <- floor(exp(0:25/10) * 1000)
> lGB <- sapply(M, function(m) {
+   st <- system.time(ret <-
+     lpmvnormR(lwr, upr, mean = mn, chol = lt, algorithm =
+       GenzBretz(maxpts = m, abseps = 0, releps = 0)))
+   return(c(st["user.self"], ll = ret))
+ })
> lH <- sapply(M, function(m) {
+   W <- NULL
+   if (require("qrng", quietly = TRUE))
+     W <- t(ghalton(m, d = J - 1))
+   st <- system.time(ret <- lpmvnorm(lwr, upr, mean = mn,
+     chol = lt, w = W, M = m))
+   return(c(st["user.self"], ll = ret))
+ })
> lHf <- sapply(M, function(m) {
+   W <- NULL
+   if (require("qrng", quietly = TRUE))
+     W <- t(ghalton(m, d = J - 1))
+   st <- system.time(ret <- lpmvnorm(lwr, upr, mean = mn, chol = lt,
+     w = W, M = m, fast = TRUE))
+   return(c(st["user.self"], ll = ret))
+ })
```

The evaluated log-likelihoods and corresponding timings are given in Figure 4.1. It seems that for  $M \geq 3000$ , results are reasonably stable.

We now define the log-likelihood function. It is important to use weights via the `w` argument (or to set the `seed`) such that only the candidate parameters `parm` change with repeated calls to `ll`. We use an extremely low number of integration points  $M$ , let's see if this still works out.

```
> M <- 500
> if (require("qrng", quietly = TRUE)) {
+   ### quasi-Monte-Carlo
+   W <- t(ghalton(M, d = J - 1))
+ } else {
+   ### Monte-Carlo
+   W <- matrix(runif(M * (J - 1)), nrow = J - 1)
+ }
> ll <- function(parm, J) {
+   m <- parm[1:J]          ### mean parameters
+   parm <- parm[-(1:J)]    ### chol parameters
+   C <- matrix(c(parm), ncol = 1L)
+   C <- ltMatrices(C, diag = TRUE, byrow = BYROW)
+   -lpmvnorm(lower = lwr, upper = upr, mean = m, chol = C,
+     w = W, M = M, logLik = TRUE)
+ }
```

We can check the correctness of our log-likelihood function




Figure 4.1: Evaluated log-likelihoods (left) and timings (right).

```
> prm <- c(mn, unclass(lt))
> ll(prm, J = J)

[1] 880.4956

> ### ATLAS gives -880.4908, M1mac gives -880.4911
> round(lpmvnormR(lwr, upr, mean = mn, chol = lt,
+               algorithm = GenzBretz(maxpts = M, abseps = 0, releps = 0)), 3)

[1] -880.491

> (llprm <- lpmvnorm(lwr, upr, mean = mn, chol = lt, w = W, M = M))

[1] -880.4956

> chk(llprm, sum(lpmvnorm(lwr, upr, mean = mn, chol = lt, w = W,
+               M = M, logLik = FALSE)))
```

Before we hand over to the optimiser, we define the score function with respect to  $\mu$  and  $\mathbf{C}$

```
> sc <- function(parm, J) {
+   m <- parm[1:J]          ### mean parameters
+   parm <- parm[-(1:J)]    ### chol parameters
+   C <- matrix(c(parm), ncol = 1L)
+   C <- ltMatrices(C, diag = TRUE, byrow = BYROW)
+   ret <- slpmvnorm(lower = lwr, upper = upr, mean = m, chol = C,
+                   w = W, M = M, logLik = TRUE)
+   return(-c(rowSums(ret$mean), rowSums(unclass(ret$chol))))
+ }
```

and check the correctness numerically

```
> if (require("numDeriv", quietly = TRUE))
+   chk(grad(ll, prm, J = J), sc(prm, J = J), check.attributes = FALSE)
```

Finally, we can hand-over to `optim`. Because we need  $\text{diag}(\mathbf{C}) > 0$ , we use box constraints and `method = "L-BFGS-B"`. We start with the estimates obtained from the original continuous data.

```
> llim <- rep(-Inf, J + J * (J + 1) / 2)
> llim[J + which(rownames(unclass(lt)) %in% paste(1:J, 1:J, sep = "."))] <- 1e-4
> if (BYROW) {
+   start <- c(rowMeans(Y), chol(Shat)[upper.tri(Shat, diag = TRUE)])
+ } else {
+   start <- c(rowMeans(Y), t(chol(Shat))[lower.tri(Shat, diag = TRUE)])
+ }
> ll(start, J = J)

[1] 875.4005

> op <- optim(start, fn = ll, gr = sc, J = J, method = "L-BFGS-B",
+           lower = llim, control = list(trace = TRUE))

iter   10 value 874.158309
final  value 874.158301
converged

> op$value ## compare with

[1] 874.1583

> ll(prm, J = J)

[1] 880.4956
```

We can now compare the true and estimated Cholesky factor  $\mathbf{C}$  of our covariance matrix  $\Sigma = \mathbf{C}\mathbf{C}^\top$

```
> (C <- ltMatrices(matrix(op$par[-(1:J)], ncol = 1),
+                 diag = TRUE, byrow = BYROW))
, , 1
      1      2      3      4
1 0.67049567 0.00000000 0.00000000 0.00000000
2 0.26764384 1.02232159 0.00000000 0.00000000
3 0.54267774 -0.05007103 1.11347760 0.00000000
4 0.05223456 0.98429745 0.08473411 0.9613685

> lt
, , 1
      1      2      3      4
1 0.7071068 0.0000000 0.0000000 0.0000000
2 0.2500000 0.9682458 0.0000000 0.0000000
3 0.6123724 -0.1581139 1.0488088 0.0000000
4 0.0000000 1.0954451 0.1651446 0.8790491
```

and the estimated means

```
> op$par[1:J]
```

```

      1      2      3      4
0.9669828 2.1281616 2.9454002 3.9886471

```

```
> mn
```

```
[1] 1 2 3 4
```

We can also compare the results on the scale of the covariance matrix

```
> ### ATLAS print issues
> round(Tcrossprod(lt), 7) ### true Sigma
```

```
, , 1
```

```

      1      2      3      4
1 0.5000000 0.1767767 0.4330127 0.000000
2 0.1767767 1.0000000 0.0000000 1.06066
3 0.4330127 0.0000000 1.5000000 0.00000
4 0.0000000 1.0606602 0.0000000 2.00000

```

```
> round(Tcrossprod(C), 7) ### interval-censored obs
```

```
, , 1
```

```

      1      2      3      4
1 0.4495644 0.1794540 0.3638631 0.0350230
2 0.1794540 1.1167747 0.0940557 1.0202488
3 0.3638631 0.0940557 1.5368386 0.0734113
4 0.0350230 1.0202488 0.0734113 1.9029791

```

```
> round(Shat, 7) ### "exact" obs
```

```

      1      2      3      4
1 0.4665566 0.1810443 0.3422224 0.0160918
2 0.1810443 0.9438534 0.0899225 0.8430953
3 0.3422224 0.0899225 1.3605492 0.0810409
4 0.0160918 0.8430953 0.0810409 1.6330153

```

This looks reasonably close.

**Warning:** Do NOT assume the choices made here (especially  $\mathbf{M}$  and  $\mathbf{W}$ ) to be universally applicable. Make sure to investigate the accuracy depending on these parameters of the log-likelihood and score function in your application.

One could ask what this whole exercise was about statistically. We estimated a multivariate normal distribution from interval-censored data, so what? Maybe we were primarily interested in fitting a linear regression

$$\mathbb{E}(Y_1 \mid Y_j = y_j, j = 2, \dots, J) = \alpha + \sum_{j=2}^J \beta_j y_j.$$

Interval-censoring in the response could have been handled by some Tobit model, but what about interval-censoring in the explanatory variables? Based on the multivariate distribution just estimated, we can obtain the regression coefficients  $\beta_j$  as

```
> c(cond_mvnorm(chol = C, which = 2:J, given = diag(J - 1))$mean)
```

```
[1] 0.2602003 0.2270392 -0.1298560
```

Alternatively, we can compute these regressions from a permuted Cholesky factor (this goes into the “simple” conditional distribution in Section 2.13)

```
> c(cond_mvnorm(chol = aperm(as.chol(C), perm = c(2:J, 1)),
+               which = 1:(J - 1), given = diag(J - 1))$mean)
```

```
[1] 0.2602003 0.2270392 -0.1298560
```

or, as a third option, from the last row of the precision matrix of the permuted Cholesky factor

```
> x <- as.array(chol2pre(aperm(as.chol(C), perm = c(2:J, 1))))[J,,1]
> c(-x[-J] / x[J])
```

```
      2      3      4
0.2602003 0.2270392 -0.1298560
```

In higher dimensions, the first option is to be preferred, because it only involves computing the Cholesky decomposition of a  $(J - 1) \times (J - 1)$  matrix, whereas the latter two options are based on a decomposition of the full  $J \times J$  covariance matrix.

We can compare these estimated regression coefficients with those obtained from a linear model fitted to the exact observations

```
> dY <- as.data.frame(t(Y))
> colnames(dY) <- paste0("Y", 1:J)
> coef(m1 <- lm(Y1 ~ ., data = dY))[-1L]
```

```
      Y2      Y3      Y4
0.3169117 0.2404565 -0.1656946
```

The estimates are quite close, but what about standard errors? Interval-censoring means loss of information, so we should see larger standard errors for the interval-censored data.

Let’s obtain the Hessian for all parameters first

```
> H <- optim(op$par, fn = ll, gr = sc, J = J, method = "L-BFGS-B",
+           lower = llim, hessian = TRUE)$hessian
```

and next we sample from the distribution of the maximum-likelihood estimators

```
> L <- try(t(chol(H)))
> ### some check on r-oldrel-macos-arm64
> if (inherits(L, "try-error"))
+   L <- t(chol(H + 1e-4 * diag(nrow(H))))
> L <- ltMatrices(L[lower.tri(L, diag = TRUE)], diag = TRUE)
> Nsim <- 50000
> Z <- matrix(rnorm(Nsim * nrow(H)), ncol = Nsim)
> rC <- solve(L, Z)[-1:J,] + op$par[-1:J] ### remove mean parameters
```

The standard error in this sample should be close to the ones obtained from the inverse Fisher information

```
> c(sqrt(rowMeans((rC - rowMeans(rC))^2)))

      5      6      7      8      9      10      11
0.05129646 0.07989618 0.12445698 0.16089554 0.07609088 0.11566519 0.14020346
      12      13      14
0.09622312 0.10415427 0.08278985

> c(sqrt(diagonals(Crossprod(solve(L)))))
```

```
[1] 0.06825507 0.10816499 0.12670329 0.14073702 0.05498052 0.10839260
[7] 0.12441885 0.14311786 0.08812684 0.11638318 0.13340466 0.09586564
[13] 0.10450821 0.08154249
```

We now coerce the matrix `rC` to an object of class `ltMatrices`

```
> rC <- ltMatrices(rC, diag = TRUE)
```

The object `rC` contains all sampled Cholesky factors of the covariance matrix. From each of these matrices, we compute the regression coefficient, giving us a sample we can use to compute standard errors from

```
> rbeta <- cond_mvnorm(chol = rC, which = 2:J, given = diag(J - 1))$mean
> sqrt(rowMeans((rbeta - rowMeans(rbeta))^2))
```

```
[1] 0.08792945 0.04869062 0.07752184
```

which are, as expected, slightly different from the ones obtained from the more informative exact observations

```
> sqrt(diag(vcov(m1)))[-1L]
```

```
          Y2          Y3          Y4
0.08229627 0.05039009 0.06246094
```

## Chapter 5

# Continuous-discrete Likelihoods

We sometimes are faced with outcomes measured at different levels of precision. Some variables might have been observed very exactly, and therefore we might want to use the log-Lebesgue density for defining the log-likelihood. Other variables might be available as relatively wide intervals only, and thus the log-likelihood is a log-probability. We can use the infrastructure developed so far to compute a joint likelihood. Let's assume we have are interested in the joint distribution of  $(\mathbf{Y}_i, \mathbf{X}_i)$  and we observed  $\mathbf{Y}_i = \mathbf{y}_i$  (that is, exact observations of  $\mathbf{Y}$ ) and  $\mathbf{a}_i < \mathbf{X}_i \leq \mathbf{b}_i$  (that is, interval-censored observations for  $\mathbf{X}_i$ ). We define the log-likelihood based on the joint normal distribution  $(\mathbf{Y}_i, \mathbf{X}_i) \sim \mathbb{N}_J((\boldsymbol{\mu}_i, \boldsymbol{\eta}_i)^\top, \mathbf{C}_i \mathbf{C}_i^\top)$  as

$$\ell_i(\boldsymbol{\mu}_i, \boldsymbol{\eta}_i, \mathbf{C}_i) = \ell_i(\boldsymbol{\mu}_i, \mathbf{C}_{\mathbf{Y},i}) + \log(\mathbb{P}(\mathbf{a}_i < \mathbf{X}_i \leq \mathbf{b}_i \mid \mathbf{C}_i, \boldsymbol{\mu}_i, \boldsymbol{\eta}_i, \mathbf{Y}_i = \mathbf{y}_i)).$$

where  $\mathbf{C}_{\mathbf{Y},i}$  is the upper part of  $\mathbf{C}_i$  corresponding to the marginal distribution of  $\mathbf{Y}_i$ . The conditional probability of  $\mathbf{X}$  given  $\mathbf{Y}$  depends on all parameters, as explained in Section 2.13. The trick here is to decompose the joint likelihood into a product of the marginal Lebesgue density of  $\mathbf{Y}_i$  and the conditional probability of  $\mathbf{X}_i$  given  $\mathbf{Y}_i = \mathbf{y}_i$ .

We first check the data

`< dp input checks 99 > ≡`

```
stopifnot(xor(missing(chol), missing(invchol)))
cJ <- nrow(obs)
dJ <- nrow(lower)
N <- ncol(obs)
stopifnot(N == ncol(lower))
stopifnot(N == ncol(upper))
if (all(mean == 0)) {
  cmean <- 0
  dmean <- 0
} else {
  if (!is.matrix(mean) || NCOL(mean) == 1L)
    mean <- matrix(mean, nrow = cJ + dJ, ncol = N)
  stopifnot(nrow(mean) == cJ + dJ)
  stopifnot(ncol(mean) == N)
  cmean <- mean[1:cJ,, drop = FALSE]
  dmean <- mean[-(1:cJ),, drop = FALSE]
}
◇
```

Fragment referenced in [100](#), [102](#).

We can use `marg_mvnorm` and `cond_mvnorm` to compute the marginal and the conditional normal distributions and the joint log-likelihood is simply the sum of the two corresponding log-likelihoods.

$\langle \text{ldpmvnorm } 100 \rangle \equiv$

```
ldpmvnorm <- function(obs, lower, upper, mean = 0, chol, invchol,
                      logLik = TRUE, ...) {

  if (missing(obs) || is.null(obs))
    return(lpmvnorm(lower = lower, upper = upper, mean = mean,
                    chol = chol, invchol = invchol, logLik = logLik, ...))
  if (missing(lower) && missing(upper) || is.null(lower) && is.null(upper))
    return(ldmvnorm(obs = obs, mean = mean,
                    chol = chol, invchol = invchol, logLik = logLik))

   $\langle \text{dp input checks } 99 \rangle$ 

  if (!missing(invchol)) {
    J <- dim(invchol)[2L]
    stopifnot(cJ + dJ == J)

    md <- marg_mvnorm(invchol = invchol, which = 1:cJ)
    ret <- ldmvnorm(obs = obs, mean = cmean, invchol = md$invchol,
                    logLik = logLik)

    cd <- cond_mvnorm(invchol = invchol, which_given = 1:cJ,
                      given = obs - cmean, center = TRUE)
    ret <- ret + lpmvnorm(lower = lower, upper = upper, mean = dmean,
                          invchol = cd$invchol, center = cd$center,
                          logLik = logLik, ...)

    return(ret)
  }

  J <- dim(chol)[2L]
  stopifnot(cJ + dJ == J)

  md <- marg_mvnorm(chol = chol, which = 1:cJ)
  ret <- ldmvnorm(obs = obs, mean = cmean, chol = md$chol, logLik = logLik)

  cd <- cond_mvnorm(chol = chol, which_given = 1:cJ,
                    given = obs - cmean, center = TRUE)
  ret <- ret + lpmvnorm(lower = lower, upper = upper, mean = dmean,
                        chol = cd$chol, center = cd$center,
                        logLik = logLik, ...)

  return(ret)
}
◇
```

Fragment referenced in [64](#).

The score function requires a little extra work. We start with the case when `invchol` is given

$\langle \text{sldpmvnorm invchol } 101 \rangle \equiv$

```

byrow_orig <- attr(invchol, "byrow")
invchol <- ltMatrices(invchol, byrow = TRUE)

J <- dim(invchol)[2L]
stopifnot(cJ + dJ == J)

md <- marg_mvnorm(invchol = invchol, which = 1:cJ)
cs <- sldmvnorm(obs = obs, mean = cmean, invchol = md$invchol)

obs_cmean <- obs - cmean
cd <- cond_mvnorm(invchol = invchol, which_given = 1:cJ,
                  given = obs_cmean, center = TRUE)
ds <- slpmvnorm(lower = lower, upper = upper, mean = dmean,
               center = cd$center, invchol = cd$invchol,
               logLik = logLik, ...)

tmp0 <- solve(cd$invchol, ds$mean, transpose = TRUE)
tmp <- - tmp0[rep(1:dJ, each = cJ),,drop = FALSE] *
      obs_cmean[rep(1:cJ, dJ),,drop = FALSE]

Jp <- nrow(unclass(invchol))
diag <- attr(invchol, "diag")
M <- as.array(ltMatrices(1:Jp, diag = diag, byrow = TRUE))[,1]
ret <- matrix(0, nrow = Jp, ncol = ncol(obs))
M1 <- M[1:cJ, 1:cJ]
idx <- t(M1)[upper.tri(M1, diag = diag)]
ret[idx,] <- Lower_tri(cs$invchol, diag = diag)

idx <- c(t(M[-(1:cJ), 1:cJ]))
ret[idx,] <- tmp

M3 <- M[-(1:cJ), -(1:cJ)]
idx <- t(M3)[upper.tri(M3, diag = diag)]
ret[idx,] <- Lower_tri(ds$invchol, diag = diag)

ret <- ltMatrices(ret, diag = diag, byrow = TRUE)
if (!diag) diagonals(ret) <- 0
ret <- ltMatrices(ret, byrow = byrow_orig)

### post differentiate mean
aL <- as.array(invchol)[-(1:cJ), 1:cJ,,drop = FALSE]
lst <- tmp0[rep(1:dJ, cJ),,drop = FALSE]
if (dim(aL)[3] == 1)
  aL <- aL[,rep(1, ncol(lst)), drop = FALSE]
dim <- dim(aL)
dobs <- -margin.table(aL * array(lst, dim = dim), 2:3)

ret <- c(list(invchol = ret, obs = cs$obs + dobs),
        ds[c("lower", "upper")])
ret$mean <- rbind(-ret$obs, ds$mean)
return(ret)

```

Fragment referenced in [102](#).

For chol, we compute the above code for its inverse and post-differentiate using the vec-trick

$\langle \text{sldpmvnorm } 102 \rangle \equiv$

```
sldpmvnorm <- function(obs, lower, upper, mean = 0, chol, invchol,
                      logLik = TRUE, ...) {

  if (missing(obs) || is.null(obs))
    return(sldpmvnorm(lower = lower, upper = upper, mean = mean,
                      chol = chol, invchol = invchol, logLik = logLik, ...))
  if (missing(lower) && missing(upper) || is.null(lower) && is.null(upper))
    return(sldpmvnorm(obs = obs, mean = mean,
                      chol = chol, invchol = invchol, logLik = logLik))

   $\langle \text{dp input checks } 99 \rangle$ 

  if (!missing(invchol)) {
     $\langle \text{sldpmvnorm invchol } 101 \rangle$ 
  }

  invchol <- solve(chol)
  ret <- sldpmvnorm(obs = obs, lower = lower, upper = upper,
                  mean = mean, invchol = invchol, logLik = logLik, ...)
  ### this means: ret$chol <- - vectrick(invchol, ret$invchol, invchol)
  ret$chol <- as.chol(- vectrick(invchol, ret$invchol))
  ret$invchol <- NULL
  return(ret)
}
◇
```

Fragment referenced in 64.

Let's assume we observed the first two dimensions exactly in our small example, and the remaining two dimensions are only known in intervals. The log-likelihood and score function for  $\mu$  and  $C$  are

```
> ic <- 1:2          ### position of continuous variables
> ll_cd <- function(parm, J) {
+   m <- parm[1:J]    ### mean parameters
+   parm <- parm[-(1:J)] ### chol parameters
+   C <- matrix(c(parm), ncol = 1L)
+   C <- ltMatrices(C, diag = TRUE, byrow = BYROW)
+   -ldpmvnorm(obs = Y[ic,], lower = lwr[-ic,],
+             upper = upr[-ic,], mean = m, chol = C,
+             w = W[-ic,,drop = FALSE], M = M)
+ }
> sc_cd <- function(parm, J) {
+   m <- parm[1:J]    ### mean parameters
+   parm <- parm[-(1:J)] ### chol parameters
+   C <- matrix(c(parm), ncol = 1L)
+   C <- ltMatrices(C, diag = TRUE, byrow = BYROW)
+   ret <- sldpmvnorm(obs = Y[ic,], lower = lwr[-ic,],
+                   upper = upr[-ic,], mean = m, chol = C,
+                   w = W[-ic,,drop = FALSE], M = M)
+   return(-c(rowSums(ret$mean),
+                   rowSums(Lower_tri(ret$chol, diag = TRUE))))
+ }
```

and the score function seems to be correct

```
> if (require("numDeriv", quietly = TRUE))
+   chk(grad(ll_cd, start, J = J), sc_cd(start, J = J),
+       check.attributes = FALSE, tol = 1e-6)
```

We can now jointly estimate all model parameters via

```
> op <- optim(start, fn = ll_cd, gr = sc_cd, J = J,
+             method = "L-BFGS-B", lower = llim,
+             control = list(trace = TRUE))
```

```
iter    10 value 655.707790
```

```
final   value 655.707779
```

```
converged
```

```
> ## estimated C
> ltMatrices(matrix(op$par[-(1:J)], ncol = 1),
+             diag = TRUE, byrow = BYROW)
```

```
, , 1
```

	1	2	3	4
1	0.68303340	0.00000000	0.00000000	0.00000000
2	0.26504369	0.93466598	0.00000000	0.00000000
3	0.53508534	-0.05736364	1.11260547	0.00000000
4	0.06748574	0.95887388	0.07774847	0.9669178

```
> ## compare with true C
```

```
> lt
```

```
, , 1
```

	1	2	3	4
1	0.7071068	0.0000000	0.0000000	0.0000000
2	0.2500000	0.9682458	0.0000000	0.0000000
3	0.6123724	-0.1581139	1.0488088	0.0000000
4	0.0000000	1.0954451	0.1651446	0.8790491

```
> ## estimated means
```

```
> op$par[1:J]
```

	1	2	3	4
	0.968533	2.126882	2.944105	3.989790

```
> ## compare with true means
```

```
> mn
```

```
[1] 1 2 3 4
```

The one restriction in both `ldpmvnorm` and `sldpmvnorm` is that the continuous variables  $\mathbf{Y}$  are ranked before the discrete variables  $\mathbf{X}$  in the observation  $(\mathbf{Y}_i, \mathbf{X}_i)$ , and thus also in  $(\boldsymbol{\mu}, \boldsymbol{\eta})$  and  $\mathbf{C}$  (the subscript  $i$  is dropped from the parameters in the following paragraph to keep the notational complexity in check).

While the means can be simply permuted, this is not the case for the Cholesky factor  $\mathbf{C}$  (see function `aperm` in Section 2.12). Of course, we can simply permute  $\hat{\mathbf{C}}_i$ , but we loose standard errors in this process. Alternatively, we can permute the order of variables in  $\mathbf{C}$  to our liking in the log-likelihood function (while keeping the original order of the observations and for the mean parameters)

```

> ### discrete variables first
> perm <- c((1:J)[-ic], ic)
> ll_ap <- function(parm, J) {
+   m <- parm[1:J]          ### mean parameters; NOT permuted
+   parm <- parm[-(1:J)]    ### chol parameters
+   C <- matrix(c(parm), ncol = 1L)
+   C <- ltMatrices(C, diag = TRUE, byrow = BYROW)
+   Ct <- aperm(as.chol(C), perm = perm)
+   -ldpmvnorm(obs = Y[ic,], lower = lwr[-ic,],
+             upper = upr[-ic,], mean = m, chol = Ct,
+             w = W[-ic,,drop = FALSE], M = M)
+ }

```

Unfortunately, this distorts the score function and we need to “de-permute” the scores. We start with  $\Sigma = \mathbf{C}\mathbf{C}^\top$ , the Cholesky decomposition of a quadratic positive definite  $J \times J$  covariance matrix. There are  $J \times (J + 1)/2$  parameters in the lower triangular part (including the diagonal) of  $\mathbf{C}$ . Changing the order of the variables by a permutation  $\pi$  with permutation matrix  $\Pi$  gives a covariance  $\Pi\mathbf{C}\mathbf{C}^\top\Pi^\top$ . This is no longer a Cholesky decomposition, because  $\Pi\mathbf{C}$  is not lower triangular. The new decomposition is

$$\Pi\mathbf{C}\mathbf{C}^\top\Pi^\top = \tilde{\mathbf{C}}\tilde{\mathbf{C}}^\top$$

( $\tilde{\mathbf{C}}$  is what `aperm` computes). As  $\mathbf{C}$ , the Cholesky factor  $\tilde{\mathbf{C}}$  is lower triangular with  $J \times (J + 1)/2$  parameters. We could write this operation as a function

$$f_3 : \mathbb{R}^{J \times (J+1)/2} \rightarrow \mathbb{R}^{J \times (J+1)/2}$$

$$f_3(\mathbf{C}) = \tilde{\mathbf{C}},$$

where in fact  $f_3 = \text{aperm}$ , and we are interested in its gradient. Deriving the gradient of a Cholesky decomposition might seem hopeless (it certainly did, at least to me, for a very long time), but there is a trick. Let us define two other functions:

$$f_1 : \mathbb{R}^{J \times (J+1)/2} \rightarrow \mathbb{R}^{J \times J}$$

$$f_1(\mathbf{C}) = \Pi\mathbf{C}\mathbf{C}^\top\Pi^\top$$

$$f_2 : \mathbb{R}^{J \times (J+1)/2} \rightarrow \mathbb{R}^{J \times J}$$

$$f_2(\tilde{\mathbf{C}}) = \tilde{\mathbf{C}}\tilde{\mathbf{C}}^\top.$$

Exploiting the chain rule for the composition  $f_1 = f_2 \circ f_3$ , we can write the gradient of  $f_1$  as the product of the gradients of  $f_2$  and  $f_3$ :

$$\frac{\partial f_1(\mathbf{C})}{\partial \mathbf{C}} = \frac{\partial f_2(\tilde{\mathbf{C}})}{\partial \tilde{\mathbf{C}}} \frac{\partial f_3(\mathbf{C})}{\partial \mathbf{C}}. \quad (5.1)$$

The last factor is what we want to compute. It turns out that it is simpler to compute the first two gradients first and, in a second step, to derive the last factor. In more detail

$$\begin{aligned} \frac{\partial f_1(\mathbf{C})}{\partial \mathbf{C}} &= \frac{\partial \Pi\mathbf{C}\mathbf{C}^\top\Pi^\top}{\partial \mathbf{C}} \\ &= \frac{\partial \Pi\mathbf{C}\mathbf{C}^\top\Pi^\top}{\partial \Pi\mathbf{C}} \frac{\partial \Pi\mathbf{C}}{\mathbf{C}} \\ &= \left( (\Pi\mathbf{C} \otimes \mathbf{I}_J) + (\mathbf{I}_J \otimes \Pi\mathbf{C}) \frac{\partial \mathbf{A}^\top}{\partial \mathbf{A}} \right) (\mathbf{I}_J \otimes \Pi). \end{aligned}$$

( $\mathbf{A}$  is a quadratic matrix and the gradient of its transpose is a permutation matrix). This analytic

expression only contains known elements and can be computed. The same applies to

$$\begin{aligned}\frac{\partial f_2(\tilde{\mathbf{C}})}{\partial \tilde{\mathbf{C}}} &= \frac{\partial \tilde{\mathbf{C}} \tilde{\mathbf{C}}^\top \Pi}{\partial \tilde{\mathbf{C}}} \\ &= (\tilde{\mathbf{C}} \otimes \mathbf{I}_J) + (\mathbf{I}_J \otimes \tilde{\mathbf{C}}) \frac{\partial \mathbf{A}^\top}{\partial \mathbf{A}}\end{aligned}$$

Both expressions treat  $\mathbf{C}$  or  $\tilde{\mathbf{C}}$  as full matrices, we are only interested in the score contributions by the  $J \times (J + 1)/2$  lower triangular elements. Using sloppy notation, we collect the relevant columns in matrices  $\mathbf{B}_1 = \frac{\partial f_1(\mathbf{C})}{\partial \mathbf{C}} \in \mathbb{R}^{J^2 \times J \times (J+1)/2}$  and  $\mathbf{B}_2 = \frac{\partial f_2(\tilde{\mathbf{C}})}{\partial \tilde{\mathbf{C}}} \in \mathbb{R}^{J^2 \times J \times (J+1)/2}$ . For the last, unknown, factor, we write  $\mathbf{B}_3 = \frac{\partial f_3(\tilde{\mathbf{C}})}{\partial \tilde{\mathbf{C}}} \in \mathbb{R}^{J \times (J+1)/2 \times J \times (J+1)/2}$  and, with formula (5.1),  $\mathbf{B}_1 = \mathbf{B}_2 \mathbf{B}_3$ . We can then solve for  $\mathbf{B}_3$  in the system  $\mathbf{B}_1^\top \mathbf{B}_1 = \mathbf{B}_1^\top \mathbf{B}_2 \mathbf{B}_3$ .

With `chol = C`, `permuted_chol = C_tilde`, `perm = pi` and score `score_schol` of the log-likelihood  $\ell(\tilde{\mathbf{C}})$  with respect to the parameters in  $\tilde{\mathbf{C}}$ , we can now implement this de-permutation of the scores. Starting with some basic sanity checks, we require lower triangular matrix objects as inputs, with diagonal elements, and check if the dimensions match

*⟨ deperma input checks chol 105a ⟩*  $\equiv$

```
stopifnot(is.ltMatrices(chol)) ### NOTE: replace with is.chol
byrow_orig <- attr(chol, "byrow")
chol <- ltMatrices(chol, byrow = FALSE)
stopifnot(is.ltMatrices(permuted_chol)) ### NOTE: replace with is.chol
permuted_chol <- ltMatrices(permuted_chol, byrow = FALSE)
stopifnot(max(abs(dim(chol) - dim(permuted_chol))) == 0)
J <- dim(chol)[2L]
stopifnot(attr(chol, "diag"))
INVCHOL <- !missing(invchol)
◇
```

Fragment referenced in 107.

Regarding `perm`, we check if this is an actual permutation

*⟨ deperma input checks perm 105b ⟩*  $\equiv$

```
if (missing(perm)) return(score_schol)
stopifnot(isTRUE(all.equal(sort(perm), 1:J)))
if (max(abs(perm - 1:J)) == 0) return(score_schol)
◇
```

Fragment referenced in 107.

The scores with respect to  $\tilde{\mathbf{C}}$  have been computed elsewhere, we just check the dimensions. In case we were given the scores with respect to  $\mathbf{L}$ , we first compute the scores with respect to  $\mathbf{C}$  (as we were lazy and only derived the results for  $\mathbf{C}$ ). As in `standardize`, the argument `score_schol` gives the score with respect to  $\mathbf{C}$  and it is the user's responsibility to provide this quantity (even when `invchol` is given).

$\langle \text{deperma input checks schol 106a} \rangle \equiv$

```

if (is.ltMatrices(score_schol)) {
  byrow_orig_s <- attr(score_schol, "byrow")
  score_schol <- ltMatrices(score_schol, byrow = FALSE)
  ### don't do this here!
  ### if (INVCHOL) score_schol <- -vectrick(permuted_invchol, score_schol)
  score_schol <- unclass(score_schol) ### this preserves byrow
}
stopifnot(is.matrix(score_schol))
N <- ncol(score_schol)
stopifnot(J * (J + 1) / 2 == nrow(score_schol))
◇

```

Fragment referenced in [107](#).

We'll have to loop over  $i = 1, \dots, N$  eventually and therefore coerce all objects to objects of class **array**, there is no need to worry about row or column storage order. We set-up indices matrices and the permutation matrix  $\Pi$

$\langle \text{deperma indices 106b} \rangle \equiv$

```

idx <- matrix(1:J^2, nrow = J, ncol = J)      ### assuming byrow = TRUE
tidx <- c(t(idx))
ltT <- idx[lower.tri(idx, diag = TRUE)]
P <- matrix(0, nrow = J, ncol = J)
P[cbind(1:J, perm)] <- 1
ID <- diag(J)
IDP <- (ID %x% P)
◇

```

Fragment referenced in [107](#).

and are now ready for the main course. We are gentle and also allow `invchol = L` as input, and we clean-up by post-differentiation at the very end in this case.

$\langle \text{deperma } 107 \rangle \equiv$

```
deperma <- function(chol = solve(invchol),
                    permuted_chol = solve(permuted_invchol),
                    invchol, permuted_invchol, perm, score_schol) {

   $\langle \text{deperma input checks chol } 105a \rangle$ 
   $\langle \text{deperma input checks perm } 105b \rangle$ 
   $\langle \text{deperma input checks schol } 106a \rangle$ 

   $\langle \text{deperma indices } 106b \rangle$ 

  Nc <- dim(chol)[1L]
  mC <- as.array(chol)[perm,,drop = FALSE]
  Ct <- as.array(permuted_chol)
  ret <- lapply(1:Nc, function(i) {
    B1 <- (mC[,i] %x% ID) + (ID %x% mC[,i])[,tidx]
    # ~~~~~ <- d t(A) / d A
    B1 <- B1 %*% IDP
    B1 <- B1[,ltT] ### relevant columns of B1
    B2 <- (Ct[,i] %x% ID) + (ID %x% Ct[,i])[,tidx]
    B2 <- B2[,ltT] ### relevant columns of B2
    B3 <- try(solve(crossprod(B2), crossprod(B2, B1)))
    if (inherits(B3, "try-error"))
      stop("failure computing permutation score")
    if (Nc == 1L)
      return(crossprod(score_schol, B3))
    return(crossprod(score_schol[,i,drop = FALSE], B3))
  })
  ret <- do.call("rbind", ret)
  ret <- ltMatrices(t(ret), diag = TRUE, byrow = FALSE)
  if (INVCHOL)
    ret <- -vectrick(chol, ret)
  ret <- ltMatrices(ret, byrow = byrow_orig_s)
  return(ret)
}
◇
```

Fragment referenced in 64.

We can now use this function to estimate the Cholesky factor for  $(\mathbf{X}, \mathbf{Y})$  when the data comes as  $(\mathbf{Y}, \mathbf{X})$  (which is needed because continuous variables come first in our implementation of log-likelihood and score function).

```
> sc_ap <- function(parm, J) {
+   m <- parm[1:J]          ### mean parameters; NOT permuted
+   parm <- parm[-(1:J)]    ### chol parameters
+   C <- matrix(c(parm), ncol = 1L)
+   C <- ltMatrices(C, diag = TRUE, byrow = BYROW)
+   ### permutation
+   Ct <- aperm(as.chol(C), perm = perm)
+   ret <- sldpmvnorm(obs = Y[ic,], lower = lwr[-ic,],
+                     upper = upr[-ic,], mean = m, chol = Ct,
+                     w = W[-ic,,drop = FALSE], M = M)
+   ### undo permutation for chol
+   retC <- deperma(chol = C, permuted_chol = Ct,
+                   perm = perm, score_schol = ret$chol)
+ }
```

```
+     return(-c(rowSums(ret$mean),
+                 rowSums(Lower_tri(retC, diag = TRUE))))
+ }
```

and the score function seems to be correct

```
> if (require("numDeriv", quietly = TRUE))
+   chk(grad(ll_ap, start, J = J), sc_ap(start, J = J),
+       check.attributes = FALSE, tol = 1e-6)
```

We can now jointly estimate all model parameters via

```
> op <- optim(start, fn = ll_ap, gr = sc_ap, J = J,
+             method = "L-BFGS-B", lower = llim,
+             control = list(trace = TRUE))
```

```
iter    10 value 655.708083
```

```
final   value 655.707779
```

```
converged
```

```
> ## estimated C for (X, Y)
> ltMatrices(matrix(op$par[-(1:J)], ncol = 1),
+             diag = TRUE, byrow = BYROW)
```

```
, , 1
```

	1	2	3	4
1	1.23595850	0.00000000	0.00000000	0.00000000
2	0.05465127	1.36451801	0.00000000	0.00000000
3	0.29575878	0.02193832	0.6152995	0.00000000
4	0.07133278	0.66705199	0.2361473	0.6618529

```
> ## compare with true _permuted_ C for (X, Y)
> aperm(as.chol(lt), perm = perm)
```

```
, , 1
```

	3	4	1	2
3	1.2247449	0.000000	0.00000000	0.000000
4	0.0000000	1.414214	0.00000000	0.000000
1	0.3535534	0.000000	0.6123724	0.000000
2	0.0000000	0.750000	0.2886751	0.595119

## Chapter 6

# Unstructured Gaussian Copula Estimation

With  $\mathbf{Z} \sim \mathbb{N}_J(0, \mathbf{I}_J)$  and  $\mathbf{Y} = \tilde{\mathbf{C}}\mathbf{Z} \sim \mathbb{N}_J(0, \tilde{\mathbf{C}}\tilde{\mathbf{C}}^\top)$  we want to estimate the off-diagonal elements of the lower triangular unit-diagonal matrix  $\mathbf{C}$ . We have  $\tilde{\mathbf{C}}(\mathbf{C}) := \text{diag}(\mathbf{C}\mathbf{C}^\top)^{-1/2}\mathbf{C}$  such that  $\boldsymbol{\Sigma} = \tilde{\mathbf{C}}\tilde{\mathbf{C}}^\top$  is a correlation matrix ( $\text{diag}(\boldsymbol{\Sigma}) = \mathbf{I}_J$ ). Note that directly estimating  $\tilde{\mathbf{C}}$  requires  $J(J+1)/2$  parameters under constraints  $\text{diag}(\boldsymbol{\Sigma}) = 1$  whereas only  $J(J-1)/2$  parameters are necessary when estimating the lower triangular part of  $\mathbf{C}$ . The standardisation by  $\text{diag}(\mathbf{C}\mathbf{C}^\top)^{-1/2}$  ensures that  $\text{diag}(\boldsymbol{\Sigma}) \equiv 1$ , that is, unconstrained optimisation can be applied.

$\langle \text{standardize } 109 \rangle \equiv$

```
standardize <- function(chol, invchol) {  
  stopifnot(xor(missing(chol), missing(invchol)))  
  if (!missing(invchol)) {  
    stopifnot(!attr(invchol, "diag"))  
    return(invcholD(invchol))  
  }  
  stopifnot(!attr(chol, "diag"))  
  return(Dchol(chol))  
}  
◇
```

Fragment referenced in [64](#).

```
> C <- ltMatrices(runif(10))  
> all.equal(as.array(chol2cov(standardize(chol = C))),  
+          as.array(chol2cor(standardize(chol = C))))  
[1] TRUE  
  
> L <- solve(C)  
> all.equal(as.array(invchol2cov(standardize(invchol = L))),  
+          as.array(invchol2cor(standardize(invchol = L))))  
[1] TRUE
```

The log-likelihood function is  $\ell_i(\mathbf{C}_i)$  (we omit  $i$  in the following) and we assume the score

$$\frac{\partial \ell(\mathbf{C})}{\partial \mathbf{C}}$$

is already available. We want to compute the score

$$\frac{\partial \ell(\tilde{\mathbf{C}})}{\partial \mathbf{C}}$$

which gives

$$\frac{\partial \ell(\tilde{\mathbf{C}})}{\partial \mathbf{C}} = \underbrace{\frac{\partial \ell(\tilde{\mathbf{C}})}{\partial \tilde{\mathbf{C}}}}_{=: \mathbf{T}} \times \frac{\partial \tilde{\mathbf{C}}(\mathbf{C})}{\partial \mathbf{C}}$$

We further have

$$\frac{\partial \tilde{\mathbf{C}}(\mathbf{C})}{\partial \mathbf{C}} = (\mathbf{C}^\top \otimes \mathbf{I}_J) \frac{\partial \text{diag}(\mathbf{C}\mathbf{C}^\top)^{-1/2}}{\partial \mathbf{C}} + (\mathbf{I}_J \otimes \text{diag}(\mathbf{C}\mathbf{C}^\top)^{-1/2})$$

and thus

$$\frac{\partial \ell(\tilde{\mathbf{C}})}{\partial \mathbf{C}} = \text{vec}(\mathbf{I}_J \mathbf{T} \mathbf{C}^\top)^\top \frac{\partial \text{diag}(\mathbf{C}\mathbf{C}^\top)^{-1/2}}{\partial \mathbf{C}} + \text{vec}(\text{diag}(\mathbf{C}\mathbf{C}^\top)^{-1/2} \mathbf{T} \mathbf{I}_J)^\top$$

and with

$$\begin{aligned} \frac{\partial \text{diag}(\mathbf{C}\mathbf{C}^\top)^{-1/2}}{\partial \mathbf{C}} &= \left. \frac{\partial \text{diag}(\mathbf{A})^{-1/2}}{\partial \mathbf{A}} \right|_{\mathbf{A}=\mathbf{C}\mathbf{C}^\top} \frac{\partial \mathbf{C}\mathbf{C}^\top}{\partial \mathbf{C}} \\ &= -\frac{1}{2} \text{diag}(\text{vec}(\text{diag}(\mathbf{C}\mathbf{C}^\top)^{-3/2})) \left[ (\mathbf{C} \otimes \mathbf{I}_J) \frac{\partial \mathbf{C}}{\partial \mathbf{C}} + (\mathbf{I}_J \otimes \mathbf{C}) \frac{\partial \mathbf{C}^\top}{\partial \mathbf{C}} \right] \end{aligned}$$

we can write

$$\begin{aligned} \text{vec}(\mathbf{I}_J \mathbf{T} \mathbf{C}^\top)^\top \left( -\frac{1}{2} \right) \text{diag}(\text{vec}(\text{diag}(\mathbf{C}\mathbf{C}^\top)^{-3/2})) &= -\frac{1}{2} \times \text{vec}(\mathbf{I}_J \mathbf{T} \mathbf{C}^\top)^\top \times \text{vec}(\text{diag}(\mathbf{C}\mathbf{C}^\top)^{-3/2})^\top \\ &=: \mathbf{b}^\top \end{aligned}$$

thus

$$\begin{aligned} \frac{\partial \ell(\tilde{\mathbf{C}})}{\partial \mathbf{C}} &= \mathbf{b}^\top \left[ (\mathbf{C} \otimes \mathbf{I}_J) \frac{\partial \mathbf{C}}{\partial \mathbf{C}} + (\mathbf{I}_J \otimes \mathbf{C}) \frac{\partial \mathbf{C}^\top}{\partial \mathbf{C}} \right] + \text{vec}(\text{diag}(\mathbf{C}\mathbf{C}^\top)^{-1/2} \mathbf{T} \mathbf{I}_J)^\top \\ &= \text{vec}(\mathbf{I}_J \mathbf{B} \mathbf{C})^\top + \text{vec}(\mathbf{C}^\top \mathbf{B} \mathbf{I}_J)^\top \frac{\partial \mathbf{C}^\top}{\partial \mathbf{C}} + \text{vec}(\text{diag}(\mathbf{C}\mathbf{C}^\top)^{-1/2} \mathbf{T} \mathbf{I}_J)^\top \end{aligned}$$

when  $\mathbf{b} = \text{vec}(\mathbf{B})$ . These scores are implemented in `destandardize` with `chol = C` and `score_schol = T`. If the model was parameterised in  $\mathbf{L} = \mathbf{C}^{-1}$ , we have `invchol = L`, however, we would still need to compute  $\mathbf{T}$  (`score_schol`, the score with respect to  $\mathbf{C}$ , and it is the user's responsibility to provide this quantity).

$\langle \text{destandardize } 111 \rangle \equiv$

```
destandardize <- function(chol = solve(invchol), invchol, score_schol)
{
  stopifnot(is.ltMatrices(chol))      ### NOTE: replace with is.chol
  J <- dim(chol)[2L]
  stopifnot(!attr(chol, "diag"))
  byrow_orig <- attr(chol, "byrow")
  chol <- ltMatrices(chol, byrow = FALSE)

  ### TODO: check byrow in score_schol?

  if (is.ltMatrices(score_schol))
    score_schol <- matrix(as.array(score_schol),
                          nrow = dim(score_schol)[2L]^2)
  stopifnot(is.matrix(score_schol))
  N <- ncol(score_schol)
  stopifnot(J^2 == nrow(score_schol))

  CCT <- Tcrossprod(chol, diag_only = TRUE)
  DC <- Dchol(chol, D = Dinv <- 1 / sqrt(CCT))
  SDC <- solve(DC)

  IDX <- t(M <- matrix(1:J^2, nrow = J, ncol = J))
  i <- cumsum(c(1, rep(J + 1, J - 1)))
  ID <- diagonals(as.integer(J), byrow = FALSE)
  if (dim(ID)[1L] != dim(chol)[1L])
    ID <- ID[rep(1, dim(chol)[1L]),]

  B <- vectrick(ID, score_schol, chol)
  B[i,] <- B[i,] * (-.5) * c(CCT)^(-3/2)
  B[-i,] <- 0

  Dtmp <- Dchol(ID, D = Dinv)

  ret <- vectrick(ID, B, chol, transpose = c(TRUE, FALSE)) +
    vectrick(chol, B, ID)[IDX,] +
    vectrick(Dtmp, score_schol, ID)

  if (!missing(invchol)) {
    ### this means: ret <- - vectrick(chol, ret, chol)
    ret <- - vectrick(chol, ret)
  }
  ret <- ret[M[lower.tri(M)],,drop = FALSE]
  if (!is.null(dimnames(chol)[[1L]]))
    colnames(ret) <- dimnames(chol)[[1L]]
  ret <- ltMatrices(ret,
                    diag = FALSE, byrow = FALSE,
                    names = dimnames(chol)[[2L]])
  ret <- ltMatrices(ret, byrow = byrow_orig)
  diagonals(ret) <- 0
  return(ret)
}
◇
```

Fragment referenced in [64](#).

We can now set-up the log-likelihood and score functions for a Gaussian copula model. We

start with the classical approach of generating the marginal observations  $\mathbf{Y}$  from the ECDF with denominator  $N + 1$  and subsequent use of the Lebesgue density as likelihood. Because no stats text on multivariate problems is complete without a reference to Edgar Anderson's iris data, let's set up a model for these four classical variables

```
> data("iris", package = "datasets")
> J <- 4
> Z <- t(qnorm(do.call("cbind", lapply(iris[1:J], rank)) / (nrow(iris) + 1)))
> (CR <- cor(t(Z)))
```

	Sepal.Length	Sepal.Width	Petal.Length	Petal.Width
Sepal.Length	1.00000000	-0.09887012	0.8695177	0.7819059
Sepal.Width	-0.09887012	1.00000000	-0.2709859	-0.2414218
Petal.Length	0.86951767	-0.27098589	1.0000000	0.8713759
Petal.Width	0.78190591	-0.24142185	0.8713759	1.0000000

```
> ll <- function(parm) {
+   C <- ltMatrices(parm)
+   Cs <- standardize(C)
+   -ldmvnorm(obs = Z, chol = Cs)
+ }
> sc <- function(parm) {
+   C <- ltMatrices(parm)
+   Cs <- standardize(C)
+   -rowSums(Lower_tri(standardize(chol = C,
+     score_schol = sldmvnorm(obs = Z, chol = Cs)$chol)))
+ }
> start <- t(chol(CR))
> start <- start[lower.tri(start)]
> if (require("numDeriv", quietly = TRUE))
+   chk(grad(ll, start), sc(start), check.attributes = FALSE)
> op <- optim(start, fn = ll, gr = sc, method = "BFGS", hessian = TRUE)
> op$value
```

```
[1] 602.5055
```

```
> S_ML <- chol2cov(standardize(ltMatrices(op$par)))
```

This approach is of course a bit strange, because we estimate the marginal distributions by nonparametric maximum likelihood whereas the joint distribution is estimated by plain maximum likelihood. For the latter, we can define the likelihood by boxes given by intervals obtained from the marginals ECDFs and estimate the Copula parameters by maximisation of this nonparametric likelihood.

```
> lwr <- do.call("cbind", lapply(iris[1:J], rank, ties.method = "min")) - 1L
> upr <- do.call("cbind", lapply(iris[1:J], rank, ties.method = "max"))
> lwr <- t(qnorm(lwr / nrow(iris)))
> upr <- t(qnorm(upr / nrow(iris)))
> M <- 500
> if (require("qrng", quietly = TRUE)) {
+   ### quasi-Monte-Carlo
+   W <- t(ghalton(M, d = J - 1))
+ } else {
+   ### Monte-Carlo
+   W <- matrix(runif(M * (J - 1)), nrow = J - 1)
```

```

+ }
> ll <- function(parm) {
+   C <- ltMatrices(parm)
+   Cs <- standardize(C)
+   -lpmvnorm(lower = lwr, upper = upr, chol = Cs, M = M, w = W)
+ }
> sc <- function(parm) {
+   C <- ltMatrices(parm)
+   Cs <- standardize(C)
+   -rowSums(Lower_tri(destandardize(chol = C,
+     score_schol = slpmvnorm(lower = lwr, upper = upr, chol = Cs,
+       M = M, w = W)$chol)))
+ }
> if (require("numDeriv", quietly = TRUE))
+   chk(grad(ll, start), sc(start), check.attributes = FALSE)
> op2 <- optim(start, fn = ll, gr = sc, method = "BFGS", hessian = TRUE)
> S_NPML <- chol2cov(standardize(ltMatrices(op2$par)))

```

For  $N = 150$ , the difference is (as expected) marginal:

```

> S_ML

, , 1

      1      2      3      4
1  1.0000000 -0.1139030  0.8768269  0.7962466
2 -0.1139030  1.0000000 -0.2856045 -0.2574850
3  0.8768269 -0.2856045  1.0000000  0.8816944
4  0.7962466 -0.2574850  0.8816944  1.0000000

> S_NPML

, , 1

      1      2      3      4
1  1.000000000 -0.09785513  0.8734599  0.7832830
2 -0.09785513  1.000000000 -0.2725997 -0.2482241
3  0.87345993 -0.27259973  1.0000000  0.8849489
4  0.78328300 -0.24822413  0.8849489  1.0000000

```

with relatively close standard errors

```

> sd_ML <- ltMatrices(sqrt(diag(solve(op$hessian))))
> diagonals(sd_ML) <- 0
> sd_NPML <- try(ltMatrices(sqrt(diag(solve(op2$hessian))))
> if (!inherits(sd_NPML, "try-error")) {
+   diagonals(sd_NPML) <- 0
+   print(sd_ML)
+   print(sd_NPML)
+ }

, , 1

      1      2      3 4
1 0.00000000 0.00000000 0.0000000 0
2 0.08122393 0.00000000 0.0000000 0

```

```

3 0.13679345 0.08761945 0.0000000 0
4 0.12621115 0.10787495 0.1010173 0

```

```

, , 1

```

```

          1          2          3 4
1 0.00000000 0.00000000 0.0000000 0
2 0.07731078 0.00000000 0.0000000 0
3 0.13999691 0.08694828 0.0000000 0
4 0.13691328 0.11037843 0.1161017 0

```

## Chapter 7

# (Experimental) User Interface

```
"interface.R" 115a≡
```

```
  < mvnrm 117a >
  < mvnrm methods 117b >
  < mvnrm simulate 118 >
  < mvnrm margDist 119 >
  < mvnrm condDist 120 >
  < mvnrm logLik 123c >
  < mvnrm LLgrad 128 >
  ◇
```

The tools provided in the previous chapters are rather low-level, so we will invest some time into setting-up a more high-level interface for representing normal models, either as  $\mathbb{N}_J(\boldsymbol{\mu}, \mathbf{C}\mathbf{C}^\top)$  or  $\mathbb{N}_J(\boldsymbol{\mu}, \mathbf{L}^{-1}\mathbf{L}^{-\top})$ , for simulating from such models, and for evaluating the log-likelihood and corresponding score functions. The latter functionality shall also work when only incomplete (variables are missing) or censored (observations are only known as intervals) data is available.

We start with the conversion of a lower triangular matrix  $\mathbf{x}$  to an `ltMatrices` object

```
< as.ltMatrices 115b > ≡
```

```
as.ltMatrices.default <- function(x) {
  stopifnot(is.numeric(x))
  if (!is.matrix(x)) x <- matrix(x)
  DIAG <- max(abs(diag(x) - 1)) > .Machine$double.eps
  DIAG <- DIAG & (nrow(x) > 1)
  lt <- x[lower.tri(x, diag = DIAG)]
  up <- x[upper.tri(x, diag = FALSE)]
  stopifnot(max(abs(up)) < .Machine$double.eps)
  nm <- rownames(x)
  if (!is.null(nm))
    return(ltMatrices(lt, diag = DIAG, names = nm))
  return(ltMatrices(lt, diag = DIAG))
}
◇
```

Fragment referenced in 2.

and proceed defining a constructor for object representing, potentially multiple, multivariate normal distributions. If the Cholesky factor  $\mathbf{C}$  (or multiple Cholesky factors  $\mathbf{C}_1, \dots, \mathbf{C}_N$ ) are

given as `chol` argument, we label them as being such objects using `as.chol`. If only a matrix is given, we convert it (if possible) to a single Cholesky factor  $\mathbf{C}$ . The same is done when  $\mathbf{L}$  is given as `invchol` argument. Of course, only one of these arguments must be specified.

$\langle mvnorm\ chol\ invchol\ 116a \rangle \equiv$

```

if (missing(chol) && missing(invchol))
  chol <- as.chol(ltMatrices(1, diag = TRUE))
stopifnot(xor(missing(chol), missing(invchol)))

if (!missing(chol)) {
  if (!is.ltMatrices(chol))
    chol <- as.ltMatrices(chol)
  scale <- as.chol(chol)
}

if (!missing(invchol)) {
  if (!is.ltMatrices(invchol))
    invchol <- as.ltMatrices(invchol)
  scale <- as.invchol(invchol)
}
ret <- list(scale = scale)
◇

```

Fragment referenced in [117a](#).

The mean, or multiple means, is stored as a  $J \times 1$  or  $J \times N$  matrix, and we check if dimensions and, possibly, names are in line with what was specified as `chol` or `invchol`

$\langle mvnorm\ mean\ 116b \rangle \equiv$

```

if (!missing(mean)) {
  stopifnot(is.numeric(mean))
  stopifnot(NROW(mean) == dim(scale)[2L])
  if (!is.matrix(mean)) {
    mean <- matrix(mean, nrow = NROW(mean))
    rownames(mean) <- names(mean)
  }
  nm <- dimnames(scale)[[2L]]
  if (is.null(rownames(mean)))
    rownames(mean) <- nm
  if (!isTRUE(all.equal(rownames(mean), nm)))
    stop("rownames of mean do not match")
  nm <- dimnames(scale)[[1L]]
  if (!is.null(nm) && dim(scale)[[2L]] == ncol(mean)) {
    if (is.null(colnames(mean)))
      colnames(mean) <- nm
    if (!isTRUE(all.equal(colnames(mean), nm)))
      stop("colnames of mean do not match")
  }
  ret$mean <- mean
}
◇

```

Fragment referenced in [117a](#).

Finally, we put everything together and return an object of class `mvnorm`, featuring `mean` and `scale`. The class of the latter slot carries the information how this object is to be interpreted (as Cholesky factor or inverse thereof)

$\langle mvnorm\ 117a \rangle \equiv$

```
### allow more than one distribution
mvnorm <- function(mean, chol, invchol) {

   $\langle mvnorm\ chol\ invchol\ 116a \rangle$ 
   $\langle mvnorm\ mean\ 116b \rangle$ 
  class(ret) <- "mvnorm"
  return(ret)
}
◇
```

Fragment referenced in [115a](#).

It might have been smarter to specify the scaled mean  $\boldsymbol{\eta} = \mathbf{L}\boldsymbol{\mu}$  because the log-density is then jointly convex in  $\boldsymbol{\eta}$  and  $\mathbf{L}$  and thus a convex problem would emerge ([Barratt and Boyd, 2023](#)).

We add a `names` and `aperm` method. The latter returns a multivariate normal distribution with permuted order of the variables

$\langle mvnorm\ methods\ 117b \rangle \equiv$

```
names.mvnorm <- function(x)
  dimnames(x$scale)[[2L]]

aperm.mvnorm <- function(a, perm, ...) {

  ret <- list(scale = aperm(a$scale, perm = perm, ...))
  if (!is.null(a$mean))
    ret$mean <- a$mean[perm,,drop = FALSE]
  class(ret) <- "mvnorm"
  ret
}
◇
```

Fragment referenced in [115a](#).

We are now ready to draw samples from such an object. If multiple normal distributions are contained in `object`, we return one sample each, otherwise, `nsim` samples are returned. Because most tools in this package expect data as  $J \times N$  matrices, we return the data in this format. If a classical `data.frame` is preferred, `as.data.frame = TRUE` we provide one

$\langle \text{mvnnorm simulate 118} \rangle \equiv$

```
simulate.mvnorm <- function(object, nsim = dim(object$scale)[1L], seed = NULL,
                             standardize = FALSE, as.data.frame = FALSE, ...) {

  J <- dim(object$scale)[2L]
  N <- dim(object$scale)[1L]
  if (N > 1)
    stopifnot(nsim == N)
  if (standardize) {
    if (is.chol(object$scale)) {
      object$scale <- standardize(chol = object$scale)
    } else {
      object$scale <- standardize(invchol = object$scale)
    }
  }
  Z <- matrix(rnorm(nsim * J), nrow = J)
  if (is.chol(object$scale)) {
    Y <- Mult(object$scale, Z)
  } else {
    Y <- solve(object$scale, Z)
  }
  ret <- Y + c(object$mean)
  rownames(ret) <- dimnames(object$scale)[[2L]]
  if (!as.data.frame)
    return(ret)
  return(as.data.frame(t(ret)))
}
◇
```

Fragment referenced in [115a](#).

It is maybe time for a first example, and we return to the iris dataset, ignoring the iris' species for the time being. We set-up a model in terms of the sample estimates

```
> data("iris", package = "datasets")
> vars <- names(iris)[-5L]
> m <- colMeans(iris[,vars])
> V <- var(iris[,vars])
> iris_mvn <- mvnnorm(mean = m, chol = t(chol(V)))
> iris_var <- simulate(iris_mvn, nsim = nrow(iris))
```

Marginal and conditional distributions might be of interest, the `margDist` and `condDist` methods are simple wrappers to `marg_mvnorm` and `cond_mvnorm`

$\langle \text{mvnorm margDist 119} \rangle \equiv$

```
margDist <- function(object, which, ...)  
  UseMethod("margDist")  
  
margDist.mvnorm <- function(object, which, ...) {  
  if (is.chol(object$scale)) {  
    ret <- list(scale = as.chol(marg_mvnorm(chol = object$scale,  
                                           which = which)$chol))  
  } else {  
    ret <- list(scale = as.invchol(marg_mvnorm(invchol = object$scale,  
                                              which = which)$invchol))  
  }  
  if (!is.null(object$mean))  
    ret$mean <- object$mean[which,,drop = FALSE]  
  class(ret) <- "mvnorm"  
  return(ret)  
}  
◇
```

Fragment referenced in [115a](#).

$\langle mvnrm\ condDist\ 120 \rangle \equiv$

```
condDist <- function(object, which_given, given, ...)
  UseMethod("condDist")

condDist.mvnrm <- function(object, which_given = 1L, given, ...) {
  if (is.chol(object$scale)) {
    ret <- cond_mvnrm(chol = object$scale, which_given = which_given,
                      given = given, ...)
    ret$scale <- as.chol(ret$chol)
    ret$chol <- NULL
  } else {
    ret <- cond_mvnrm(invchol = object$scale, which_given = which_given,
                      given = given, ...)
    ret$invchol <- as.chol(ret$invchol)
    ret$invchol <- NULL
  }
  if (!is.null(object$mean)) {
    if (is.character(which_given))
      which_given <- match(which_given, dimnames(object$scale)[[2L]])
    if (ncol(object$mean) > 1L && ncol(ret$mean) > 1)
      stop("dimensions do not match")
    if (ncol(object$mean) == 1L && ncol(ret$mean) > 1L) {
      ret$mean <- object$mean[-which_given,,drop = TRUE] + ret$mean
    } else {
      ret$mean <- object$mean[-which_given,,drop = FALSE] + c(ret$mean)
    }
  }
  class(ret) <- "mvnrm"
  return(ret)
}
◇
```

Fragment referenced in [115a](#).

We could now compute the marginal distribution of two Petal variables or the bivariate regressions of the two Petal variables given the observed Sepal variables. Note that the last object contains  $N = 150$  different distributions

```
> j <- 3:4
> margDist(iris_mvn, which = vars[j])

$scale
, , 1

      Petal.Length Petal.Width
Petal.Length  1.7652982  0.0000000
Petal.Width   0.7339323  0.2057903

$mean
      [,1]
Petal.Length 3.758000
Petal.Width  1.199333
```

```
attr("class")
[1] "mvnorm"

> gm <- t(iris[,vars[-(j)]])
> iris_cmvn <- condDist(iris_mvn, which = vars[j], given = gm)
```

We now work towards implementating the corresponding log-likelihood function. This is a trivial task as long as all variables for all observations have been observed exactly (that is, we can interpret the data as being continuous). Here, we also want to allow imprecise, that is, interval-censored, measurements. The one constraint in `ldpmvnorm` is that the continuous variables come first, followed by the censored ones. This of course might not be in line with the variable ordering we have in mind for our model. Our log-likelihood function shall be able to evaluate the log-likelihood for arbitrary permutations of the variables and, optionally, also based on marginal distributions in case observations are missing.

The following `logLik` method for objects of class `mvnorm` is essentially a wrapper for `ldpmvnorm`, handling permutations, marginalisation, and standardisation. We begin with some sanity checks

$\langle \text{argchecks } 122 \rangle \equiv$

```
args <- c(object, list(...))
nargs <- missing(obs) + missing(lower) + missing(upper)
stopifnot(nargs < 3L)

nmobs <- NULL
if (!missing(obs)) {
  if (!is.null(obs)) {
    stopifnot(is.matrix(obs))
    nmobs <- rownames(obs)
  }
}
nmlower <- nmupper <- nmlu <- NULL
if (!missing(lower)) {
  if (!is.null(lower)) {
    stopifnot(is.matrix(lower))
    nmlu <- nmlower <- rownames(lower)
  }
}
if (!missing(upper)) {
  if (!is.null(lower)) {
    stopifnot(is.matrix(upper))
    nmupper <- rownames(upper)
    if (!missing(lower)) {
      stopifnot(isTRUE(all.equal(nmlower, nmupper)))
    } else {
      nmlu <- nmupper
    }
  }
}

nm <- c(nmobs, nmlu)
no <- names(object)
stopifnot(nm %in% no)
perm <- NULL
if (!isTRUE(all.equal(nm, no)))
  perm <- c(nm, no[!no %in% nm])

if (!missing(obs)) args$obs <- obs
if (!missing(lower)) args$lower <- lower
if (!missing(upper)) args$upper <- upper
◇
```

Fragment referenced in [123c](#), [128](#).

and proceed with the workhorse when **C** was given

$\langle \logLik\ chol\ 123a \rangle \equiv$

```
names(args)[names(args) == "scale"] <- "chol"
if (standardize)
  args$chol <- standardize(chol = args$chol)
if (!is.null(perm)) {
  args$chol <- aperm(as.chol(args$chol), perm = perm)
  if (length(nm) < length(no))
    args$chol <- marg_mvnorm(chol = args$chol, which = nm)$chol
  args$mean <- args$mean[nm,,drop = FALSE]
}
return(do.call("ldpmvnorm", args))
◇
```

Fragment referenced in [123c](#).

For inverse Cholesky factors **L**, the code is very similar, just the argument names change

$\langle \logLik\ invchol\ 123b \rangle \equiv$

```
names(args)[names(args) == "scale"] <- "invchol"
if (standardize)
  args$invchol <- standardize(invchol = args$invchol)
if (!is.null(perm)) {
  args$invchol <- aperm(as.invchol(args$invchol), perm = perm)
  if (length(nm) < length(no))
    args$invchol <- marg_mvnorm(invchol = args$invchol,
                                which = nm)$invchol
  args$mean <- args$mean[nm,,drop = FALSE]
}
return(do.call("ldpmvnorm", args))
◇
```

Fragment referenced in [123c](#).

Putting everything together in a corresponding `logLik` method

$\langle mvnorm\ logLik\ 123c \rangle \equiv$

```
logLik.mvnorm <- function(object, obs, lower, upper, standardize = FALSE,
  ...) {
   $\langle argchecks\ 122 \rangle$ 
  if (is.chol(object$scale)) {
     $\langle logLik\ chol\ 123a \rangle$ 
  }
   $\langle logLik\ invchol\ 123b \rangle$ 
}
◇
```

Fragment referenced in [115a](#).

allows us to evaluate the log-likelihood of the conditional models for iris

```
> logLik(object = iris_cmvn, obs = t(iris[,vars[-j]]))
```

```
[1] -4750.456
```

This implementation of the log-likelihood silently handles the case when variables have been specified in a different order than hard-wired into the model

```
> logLik(object = iris_cmvn, obs = t(iris[,rev(vars[-j])]))

[1] -4750.456
```

The hardest task is the implementation of a score function which features the same options as the log-likelihood function and provides the gradients with respect not only to the parameters ( $\mu$  and **C** or **L**), but also with respect to the data objects **obs**, **lower**, and **upper**.

In essence, we have to repair the damage imposed by a series of transformations in `logLik.mvnorm`, that is, by standardisation, permutation, and marginalisation. We start with the case when **C** was given. First, we repeat all the steps performed in `logLik`, but call the score function `sldpmvnorm` instead of the log-likelihood function `ldpmvnorm`

```
< lLgrad chol 124a > ≡

names(args)[names(args) == "scale"] <- "chol"
sc <- args$chol
if (standardize)
  args$chol <- sc <- standardize(chol = args$chol)
if (!is.null(perm)) {
  if (!attr(args$chol, "diag")) {
    diagonals(args$chol) <- 1
    sc <- args$chol
  }
  args$chol <- pc <- aperm(as.chol(args$chol), perm = perm)
  if (length(nm) < length(no))
    args$chol <- marg_mvnorm(chol = args$chol, which = nm)$chol
  args$mean <- args$mean[nm,,drop = FALSE]
}
ret <- do.call("sldpmvnorm", args)
< lLgrad mean 124b >
< lLgrad marginalisation 125a >
< lLgrad deperma 125b >
< lLgrad destandarized 125c >
< lLgrad diagonals 126a >
< lLgrad return 126b >
◇
```

Fragment referenced in 128.

The next task is to post-differentiate all scores such that the gradients with respect to the original arguments of `logLik` are obtained. We start with the gradient with respect to  $\mu$ , in case it was not given

```
< lLgrad mean 124b > ≡

### sldmvnorm returns mean score as -obs
if (is.null(ret$mean)) ret$mean <- - ret$obs
◇
```

Fragment referenced in 124a.

In case we marginalised over some variables, we have to set the omitted parameters to zero

$\langle \text{LLgrad marginalisation 125a} \rangle \equiv$

```
om <- length(no) - length(nm)
if (om > 0) {
  am <- matrix(0, nrow = om, ncol = ncol(ret$mean))
  rownames(am) <- no[!no %in% nm]
  ret$mean <- rbind(ret$mean, am)
  Jo <- dim(object$scale)[[2L]]
  pJ <- dim(args$invchol)[[2L]]
  am <- matrix(0, nrow = Jo * (Jo + 1) / 2 - pJ * (pJ + 1) / 2,
               ncol = dim(ret$invchol)[1L])
  byrow_orig <- attr(ret$chol, "byrow")
  ret$chol <- ltMatrices(ret$chol, byrow = TRUE)
  ### rbind only works for byrow = TRUE
  ret$chol <- ltMatrices(rbind(unclass(ret$chol), am),
                        byrow = TRUE,
                        diag = TRUE,
                        names = perm)
  ret$chol <- ltMatrices(ret$chol, byrow = byrow_orig)
}
◇
```

Fragment referenced in [124a](#).

If the order of the variables was permuted, we compute the scores for the original ordering of the variables, as explained in [Chapter 5](#)

$\langle \text{LLgrad deperma 125b} \rangle \equiv$

```
if (!is.null(perm))
  ret$chol <- deperma(chol = sc, permuted_chol = pc,
                    perm = match(perm, no),
                    score_schol = ret$chol)
◇
```

Fragment referenced in [124a](#).

The effect of standardization can be removed as discussed in [Chapter 6](#)

$\langle \text{LLgrad destandarized 125c} \rangle \equiv$

```
if (standardize)
  ret$chol <- destandardize(chol = object$scale,
                          score_schol = ret$chol)
◇
```

Fragment referenced in [124a](#).

and it remains to remove fix diagonal elements

$\langle \text{LLgrad diagonals 126a} \rangle \equiv$

```

if (!attr(sc, "diag"))
  ret$chol <- ltMatrices(Lower_tri(ret$chol, diag = FALSE),
                        diag = FALSE,
                        byrow = attr(ret$chol, "byrow"),
                        names = dimnames(ret$chol)[[2L]])

```

◇

Fragment referenced in [124a](#).

and to return the results, with mean scores in the correct ordering

$\langle \text{LLgrad return 126b} \rangle \equiv$

```

ret$scale <- ret$chol
ret$chol <- NULL
ret$mean <- ret$mean[no,,drop = FALSE]
return(ret)

```

◇

Fragment referenced in [124a](#).

The steps are essentially the same when  $\mathbf{L}$  was given, but we have to post-differentiate  $\mathbf{C} = \mathbf{L}^{-1}$  with respect to  $\mathbf{L}$

$\langle \text{Lgrad invchol } 127 \rangle \equiv$

```

names(args)[names(args) == "scale"] <- "invchol"
si <- args$invchol
if (standardize)
  args$invchol <- si <- standardize(invchol = args$invchol)
if (!is.null(perm)) {
  if (!attr(args$invchol, "diag")) {
    diagonals(args$invchol) <- 1
    si <- args$invchol
  }
  args$invchol <- pi <- aperm(as.invchol(args$invchol), perm = perm)
  if (length(nm) < length(no))
    args$invchol <- marg_mvnorm(invchol = args$invchol,
                               which = nm)$invchol
  args$mean <- args$mean[nm,,drop = FALSE]
}
ret <- do.call("sldpmvnorm", args)
### sldmvnorm returns mean score as -obs
if (is.null(ret$mean)) ret$mean <- - ret$obs
om <- length(no) - length(nm)
if (om > 0) {
  am <- matrix(0, nrow = om, ncol = ncol(ret$mean))
  rownames(am) <- no[!no %in% nm]
  ret$mean <- rbind(ret$mean, am)
  Jo <- dim(object$scale)[[2L]]
  pJ <- dim(args$invchol)[[2L]]
  am <- matrix(0, nrow = Jo * (Jo + 1) / 2 - pJ * (pJ + 1) / 2,
              ncol = dim(ret$invchol)[1L])
  byrow_orig <- attr(ret$invchol, "byrow")
  ret$invchol <- ltMatrices(ret$invchol, byrow = TRUE)
  ### rbind only works for byrow = TRUE
  ret$invchol <- ltMatrices(rbind(unclass(ret$invchol), am),
                           byrow = TRUE,
                           diag = TRUE,
                           names = perm)
  ret$invchol <- ltMatrices(ret$invchol, byrow = byrow_orig)
}
if (!is.null(perm))
  ret$invchol <- deperma(invchol = si, permuted_invchol = pi,
                       perm = match(perm, no),
                       score_schol = -vectrick(pi, ret$invchol))
if (standardize)
  ret$invchol <- destandardize(invchol = object$scale,
                              score_schol = -vectrick(si, ret$invchol))
if (!attr(si, "diag"))
  ret$invchol <- ltMatrices(Lower_tri(ret$invchol, diag = FALSE),
                           diag = FALSE,
                           byrow = attr(ret$invchol, "byrow"),
                           names = dimnames(ret$invchol)[[2L]])

ret$scale <- ret$invchol
ret$invchol <- NULL
ret$mean <- ret$mean[no,,drop = FALSE]
return(ret)

```

Fragment referenced in [128](#).

We can now provide the log-likelihood gradients

$\langle \text{mvnorm } \text{llgrad } 128 \rangle \equiv$

```
llgrad <- function(object, ...)
  UseMethod("llgrad")

llgrad.mvnorm <- function(object, obs, lower, upper, standardize = FALSE,
  ...) {
   $\langle \text{argchecks } 122 \rangle$ 
  if (is.chol(object$scale)) {
     $\langle \text{llgrad chol } 124a \rangle$ 
  }
   $\langle \text{llgrad invchol } 127 \rangle$ 
}
 $\diamond$ 
```

Fragment referenced in [115a](#).

Let's use this infrastructure to set-up maximum-likelihood estimation procedures. We start implementing the log-likelihood and score functions for the iris dataset

```
> J <- length(vars)
> obs <- t(iris[, vars])
> lower <- upper <- NULL
> ll <- function(parm) {
+   C <- ltMatrices(parm[-(1:J)], diag = TRUE, names = vars)
+   x <- mvnorm(mean = parm[1:J], chol = C)
+   -logLik(object = x, obs = obs, lower = lower, upper = upper)
+ }
> sc <- function(parm) {
+   C <- ltMatrices(parm[-(1:J)], diag = TRUE, names = vars)
+   x <- mvnorm(mean = parm[1:J], chol = C)
+   ret <- llgrad(object = x, obs = obs, lower = lower, upper = upper)
+   -c(rowSums(ret$mean), rowSums(Lower_tri(ret$scale, diag = TRUE)))
+ }
```

and can now estimate the mean and Cholesky factor of the covariance matrix

```
> start <- c(c(iris_mvn$mean), Lower_tri(iris_mvn$scale, diag = TRUE))
> if (require("numDeriv", quietly = TRUE))
+   chk(grad(ll, start), sc(start), check.attributes = FALSE)
> op <- optim(start, fn = ll, gr = sc, method = "L-BFGS-B",
+   lower = llim, control = list(trace = TRUE))

iter   10 value 379.914996
iter   20 value 379.914642
final   value 379.914631
converged

> Chat <- ltMatrices(op$par[-(1:J)], diag = TRUE, names = vars)
> ML <- mvnorm(mean = op$par[1:J], chol = Chat)
```

Quit unsurprisingly, the results are practically equivalent to the analytically available maximum-likelihood estimators in this case

```

> ### covariance
> round(chol2cov(ML$scale), 2)

, , 1

      Sepal.Length Sepal.Width Petal.Length Petal.Width
Sepal.Length      0.68      -0.04        1.27        0.51
Sepal.Width       -0.04       0.19       -0.33       -0.12
Petal.Length      1.27      -0.33        3.10        1.29
Petal.Width       0.51      -0.12        1.29        0.58

> N <- nrow(iris)
> round(V * (N - 1) / N, 2)

      Sepal.Length Sepal.Width Petal.Length Petal.Width
Sepal.Length      0.68      -0.04        1.27        0.51
Sepal.Width       -0.04       0.19       -0.33       -0.12
Petal.Length      1.27      -0.33        3.10        1.29
Petal.Width       0.51      -0.12        1.29        0.58

> ### mean
> ML$mean[, , drop = TRUE]

Sepal.Length Sepal.Width Petal.Length Petal.Width
5.843333     3.057333     3.758000     1.199333

> m

Sepal.Length Sepal.Width Petal.Length Petal.Width
5.843333     3.057333     3.758000     1.199333

```

Now, this was a lot of work to replace `mean` and `var` with something more fancy, and we would of course not go down this way in real life. But how about a more complex situation where one (or more) variables are only known up to intervals? Let's present the first variable is such a case

```

> v1 <- vars[1]
> q1 <- quantile(iris[[v1]], prob = 1:4 / 5)
> head(f1 <- cut(iris[[v1]], breaks = c(-Inf, q1, Inf)))

[1] (5,5.6] (-Inf,5] (-Inf,5] (-Inf,5] (-Inf,5] (5,5.6]
Levels: (-Inf,5] (5,5.6] (5.6,6.1] (6.1,6.52] (6.52, Inf]

```

The only necessary modification to our code is the specification of `lower` and `upper` bounds for these intervals, and the removal of the first variable from the “exact continuous” observations `obs`. The rest of the machinery *doesn't need any update at all*. Note that the mean and covariance parameters are no longer orthogonal (as in the toy example above), so we do have to optimise over both sets of parameters simultaneously.

```

> lower <- matrix(c(-Inf, q1)[f1], nrow = 1)
> upper <- matrix(c(q1, Inf)[f1], nrow = 1)
> rownames(lower) <- rownames(upper) <- v1
> obs <- obs[!rownames(obs) %in% v1, , drop = FALSE]
> if (require("numDeriv", quietly = TRUE))
+   chk(grad(ll, start), sc(start), check.attributes = FALSE)
> opi <- optim(start, fn = ll, gr = sc, method = "L-BFGS-B",
+             lower = llim, control = list(trace = TRUE))

```

```

iter    10 value 472.264901
iter    20 value 472.211584
iter    30 value 472.208646
iter    40 value 472.208420
final   value 472.208416
converged

```

```

> Chati <- ltMatrices(opi$par[-(1:J)], diag = TRUE, names = vars)
> MLi <- mvnrm(mean = opi$par[1:J], chol = Chati)

```

Because the likelihood is a product of a continuous density and a conditional probability as introduced in Chapter 5, the two in-sample log-likelihoods are not comparable. However, the parameters of the two estimated normal distributions can be compared directly (and are rather close in our case)

```
> op$value
```

```
[1] 379.9146
```

```
> opi$value
```

```
[1] 472.2084
```

```
> ### covariance
```

```
> round(chol2cov(MLi$scale), 2)
```

```
, , 1
```

	Sepal.Length	Sepal.Width	Petal.Length	Petal.Width
Sepal.Length	0.73	-0.03	1.27	0.52
Sepal.Width	-0.03	0.19	-0.33	-0.12
Petal.Length	1.27	-0.33	3.09	1.29
Petal.Width	0.52	-0.12	1.29	0.58

```
> round(chol2cov(ML$scale), 2)
```

```
, , 1
```

	Sepal.Length	Sepal.Width	Petal.Length	Petal.Width
Sepal.Length	0.68	-0.04	1.27	0.51
Sepal.Width	-0.04	0.19	-0.33	-0.12
Petal.Length	1.27	-0.33	3.10	1.29
Petal.Width	0.51	-0.12	1.29	0.58

```
> ### mean
```

```
> MLi$mean[,drop = TRUE]
```

Sepal.Length	Sepal.Width	Petal.Length	Petal.Width
5.759601	3.057314	3.757959	1.199314

```
> ML$mean[,drop = TRUE]
```

Sepal.Length	Sepal.Width	Petal.Length	Petal.Width
5.843333	3.057333	3.758000	1.199333

We close this chapter with a word of warning: If more than one variable is censored, the `M` and `w` arguments to `lpmvnorm` and `slpmvnorm` have to be specified in `logLik` and `lLgrad` as additional arguments (...) *AND MUST BE IDENTICAL* in both calls.

## Chapter 8

# Package Infrastructure

$\langle R \text{ Header } 131 \rangle \equiv$

```
###      Copyright (C) 2022- Torsten Hothorn
###
###      This file is part of the 'mvtnorm' R add-on package.
###
###      'mvtnorm' is free software: you can redistribute it and/or modify
###      it under the terms of the GNU General Public License as published by
###      the Free Software Foundation, version 2.
###
###      'mvtnorm' is distributed in the hope that it will be useful,
###      but WITHOUT ANY WARRANTY; without even the implied warranty of
###      MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
###      GNU General Public License for more details.
###
###      You should have received a copy of the GNU General Public License
###      along with 'mvtnorm'. If not, see <http://www.gnu.org/licenses/>.
###
###
###      DO NOT EDIT THIS FILE
###
###      Edit 'lmvnorm_src.w' and run 'nuweb -r lmvnorm_src.w'
```

◇

Fragment referenced in [2](#), [64](#).

$\langle C \text{ Header } 132 \rangle \equiv$

```
/*
  Copyright (C) 2022- Torsten Hothorn

  This file is part of the 'mvtnorm' R add-on package.

  'mvtnorm' is free software: you can redistribute it and/or modify
  it under the terms of the GNU General Public License as published by
  the Free Software Foundation, version 2.

  'mvtnorm' is distributed in the hope that it will be useful,
  but WITHOUT ANY WARRANTY; without even the implied warranty of
  MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
  GNU General Public License for more details.

  You should have received a copy of the GNU General Public License
  along with 'mvtnorm'. If not, see <http://www.gnu.org/licenses/>.

  DO NOT EDIT THIS FILE

  Edit 'lmvnorm_src.w' and run 'nuweb -r lmvnorm_src.w'
*/
◇
```

Fragment referenced in [3](#), [65](#).

# Appendix

This document uses the following matrix derivatives

$$\begin{aligned}
\frac{\partial \mathbf{y}^\top \mathbf{A}^\top \mathbf{A} \mathbf{y}}{\partial \mathbf{A}} &= 2\mathbf{A} \mathbf{y} \mathbf{y}^\top \\
\frac{\partial \mathbf{A}^{-1}}{\partial \mathbf{A}} &= -(\mathbf{A}^{-\top} \otimes \mathbf{A}^{-1}) \\
\frac{\partial \mathbf{A} \mathbf{A}^\top}{\partial \mathbf{A}} &= (\mathbf{A} \otimes \mathbf{I}_J) \frac{\partial \mathbf{A}}{\partial \mathbf{A}} + (\mathbf{I}_J \otimes \mathbf{A}) \frac{\partial \mathbf{A}^\top}{\partial \mathbf{A}} \\
&= (\mathbf{A} \otimes \mathbf{I}_J) + (\mathbf{I}_J \otimes \mathbf{A}) \frac{\partial \mathbf{A}^\top}{\partial \mathbf{A}} \\
\frac{\partial \text{diag}(\mathbf{A})}{\partial \mathbf{A}} &= \text{diag}(\text{vec}(\mathbf{I}_J)) \\
\frac{\partial \mathbf{A}}{\partial \mathbf{A}} &= \text{diag}(I_{J^2}) \\
\frac{\partial \mathbf{y}^\top \mathbf{A} \mathbf{y}}{\partial \mathbf{y}} &= \mathbf{y}^\top (\mathbf{A} + \mathbf{A}^\top) \\
\frac{\partial \mathbf{B} \mathbf{A}}{\partial \mathbf{A}} &= (\mathbf{I}_J \otimes \mathbf{B})
\end{aligned}$$

and the “vec trick”  $\text{vec}(\mathbf{X})^\top (\mathbf{B} \otimes \mathbf{A}^\top) = \text{vec}(\mathbf{A} \mathbf{X} \mathbf{B})^\top$ .

# Index

## Files

"interface.R" Defined by [115a](#).  
"lpmvnorm.c" Defined by [65](#).  
"lpmvnorm.R" Defined by [64](#).  
"ltMatrices.c" Defined by [3](#).  
"ltMatrices.R" Defined by [2](#).

## Fragments

[〈.subset ltMatrices 13〉](#) Referenced in [14](#).  
[〈add diagonal elements 20〉](#) Referenced in [2](#).  
[〈aperm 51ab〉](#) Referenced in [2](#).  
[〈aperm checks 50〉](#) Referenced in [51a](#).  
[〈argchecks 122〉](#) Referenced in [123c](#), [128](#).  
[〈as.ltMatrices 115b〉](#) Referenced in [2](#).  
[〈assign diagonal elements 21〉](#) Referenced in [2](#).  
[〈C Header 132〉](#) Referenced in [3](#), [65](#).  
[〈C length 24a〉](#) Referenced in [24b](#), [26](#), [29](#), [33a](#), [42a](#).  
[〈check A argument 43b〉](#) Referenced in [44](#).  
[〈check and / or set integration weights 74b〉](#) Referenced in [75](#), [87](#).  
[〈check C argument 42b〉](#) Referenced in [44](#).  
[〈check obs 57b〉](#) Referenced in [2](#).  
[〈check S argument 43a〉](#) Referenced in [44](#).  
[〈chol 40〉](#) Referenced in [3](#).  
[〈chol classes 45〉](#) Referenced in [48](#).  
[〈chol scores 77a〉](#) Referenced in [78a](#).  
[〈chol syMatrices 39〉](#) Referenced in [2](#).  
[〈Cholesky of precision 74c〉](#) Referenced in [75](#), [87](#).  
[〈colSumsdnorm 58a〉](#) Referenced in [3](#).  
[〈colSumsdnorm ltMatrices 58b〉](#) Referenced in [2](#).  
[〈compute x 68b〉](#) Referenced in [69b](#), [83a](#).  
[〈compute y 68a〉](#) Referenced in [69b](#), [83a](#).  
[〈cond general 53〉](#) Referenced in [55](#).  
[〈cond simple 54〉](#) Referenced in [55](#).  
[〈conditional 55〉](#) Referenced in [2](#).  
[〈convenience functions 48〉](#) Referenced in [2](#).  
[〈crossprod ltMatrices 38〉](#) Referenced in [2](#).  
[〈D times C 46〉](#) Referenced in [48](#).  
[〈deperma 107〉](#) Referenced in [64](#).  
[〈deperma indices 106b〉](#) Referenced in [107](#).  
[〈deperma input checks chol 105a〉](#) Referenced in [107](#).  
[〈deperma input checks perm 105b〉](#) Referenced in [107](#).  
[〈deperma input checks schol 106a〉](#) Referenced in [107](#).  
[〈destandardize 111〉](#) Referenced in [64](#).

<diagonal matrix 22> Referenced in 2.  
 <diagonals ltMatrices 19> Referenced in 2.  
 <dim ltMatrices 6c> Referenced in 2.  
 <dimensions 71c> Referenced in 73, 84.  
 <dimnames ltMatrices 7a> Referenced in 2.  
 <dp input checks 99> Referenced in 100, 102.  
 <extract slots 10> Referenced in 11, 12, 13, 17, 19, 21, 23a, 27.  
 <first element 34a> Referenced in 34b, 35a.  
 <IDX 35b> Referenced in 36, 42a.  
 <increment 69c> Referenced in 73.  
 <init center 72c> Referenced in 73, 84.  
 <init dans 83c> Referenced in 84.  
 <init logLik loop 67c> Referenced in 73, 79b.  
 <init random seed, reset on exit 74a> Referenced in 75, 87.  
 <init score loop 79b> Referenced in 84.  
 <initialisation 67b> Referenced in 73, 84.  
 <inner logLik loop 69b> Referenced in 73.  
 <inner score loop 83a> Referenced in 84.  
 <input checks 66> Referenced in 63, 75, 87.  
 <is.ltMatrices 7c> Referenced in 2.  
 <kronecker vec trick 44> Referenced in 2.  
 <L times D 47> Referenced in 48.  
 <lapack options 28> Referenced in 29, 30.  
 <ldmvnorm 57a> Referenced in 64.  
 <ldmvnorm chol 59a> Referenced in 57a.  
 <ldmvnorm invchol 59b> Referenced in 57a.  
 <ldpmvnorm 100> Referenced in 64.  
 <llgrad chol 124a> Referenced in 128.  
 <llgrad deperma 125b> Referenced in 124a.  
 <llgrad destandardized 125c> Referenced in 124a.  
 <llgrad diagonals 126a> Referenced in 124a.  
 <llgrad invchol 127> Referenced in 128.  
 <llgrad marginalisation 125a> Referenced in 124a.  
 <llgrad mean 124b> Referenced in 124a.  
 <llgrad return 126b> Referenced in 124a.  
 <logdet 33a> Referenced in 3.  
 <logdet ltMatrices 33b> Referenced in 2.  
 <logLik chol 123a> Referenced in 123c.  
 <logLik invchol 123b> Referenced in 123c.  
 <lower scores 77c> Referenced in 78a.  
 <lower triangular elements 17> Referenced in 2.  
 <lpmvnorm 75> Referenced in 64.  
 <lpmvnormR 63> Not referenced.  
 <ltMatrices 6a> Referenced in 2.  
 <ltMatrices dim 4> Referenced in 6a.  
 <ltMatrices input 5b> Referenced in 6a.  
 <ltMatrices names 5a> Referenced in 6a.  
 <marginal 52b> Referenced in 2.  
 <mc input checks 52a> Referenced in 52b, 55.  
 <mean scores 77b> Referenced in 78a.  
 <move on 70a> Referenced in 73, 84.  
 <mult 24b> Referenced in 3.  
 <mult ltMatrices 23a> Referenced in 2.  
 <mult ltMatrices transpose 25> Referenced in 23a.  
 <mult syMatrices 27> Referenced in 2.  
 <mult transpose 26> Referenced in 3.  
 <mvnorm 117a> Referenced in 115a.  
 <mvnorm chol invchol 116a> Referenced in 117a.

<mvnorm condDist 120> Referenced in 115a.  
 <mvnorm lLgrad 128> Referenced in 115a.  
 <mvnorm logLik 123c> Referenced in 115a.  
 <mvnorm margDist 119> Referenced in 115a.  
 <mvnorm mean 116b> Referenced in 117a.  
 <mvnorm methods 117b> Referenced in 115a.  
 <mvnorm simulate 118> Referenced in 115a.  
 <names ltMatrices 7b> Referenced in 2.  
 <new score means, lower and upper 81c> Referenced in 83a.  
 <output 69d> Referenced in 73.  
 <pnorm 71a> Referenced in 73, 84.  
 <pnorm fast 70b> Referenced in 65.  
 <pnorm slow 70c> Referenced in 65.  
 <post differentiate chol score 85d> Referenced in 87.  
 <post differentiate invchol score 86a> Referenced in 87.  
 <post differentiate lower score 85b> Referenced in 87.  
 <post differentiate mean score 85a> Referenced in 87.  
 <post differentiate upper score 85c> Referenced in 87.  
 <post process score 86b> Referenced in 87.  
 <print ltMatrices 11> Referenced in 2.  
 <R Header 131> Referenced in 2, 64.  
 <R lpmvnorm 73> Referenced in 65.  
 <R slpmvnorm 84> Referenced in 65.  
 <R slpmvnorm variables 72d> Referenced in 73, 84.  
 <RC input 23b> Referenced in 24b, 26, 29, 30, 33a, 36, 42a.  
 <reorder ltMatrices 12> Referenced in 2.  
 <score a, b 79a> Referenced in 79b, 84.  
 <score c11 78b> Referenced in 79b, 84.  
 <score output 83b> Referenced in 84.  
 <score output object 78a> Referenced in 84.  
 <score wrt new chol diagonal 81b> Referenced in 83a.  
 <score wrt new chol off-diagonals 81a> Referenced in 83a.  
 <setup return object 72a> Referenced in 73.  
 <sldmvnorm 61> Referenced in 64.  
 <sldpmvnorm 102> Referenced in 64.  
 <sldpmvnorm invchol 101> Referenced in 102.  
 <slpmvnorm 87> Referenced in 64.  
 <solve 29> Referenced in 3.  
 <solve C 30> Referenced in 3.  
 <solve ltMatrices 31> Referenced in 2.  
 <standardise 67a> Referenced in 75, 87.  
 <standardize 109> Referenced in 64.  
 <subset ltMatrices 14> Referenced in 2.  
 <syMatrices 6b> Referenced in 2.  
 <t(C) S t(A) 41> Referenced in 42a.  
 <tcrossprod 36> Referenced in 3.  
 <tcrossprod diagonal only 34b> Referenced in 36.  
 <tcrossprod full 35a> Referenced in 36.  
 <tcrossprod ltMatrices 37> Referenced in 2.  
 <univariate problem 72b> Referenced in 73.  
 <update d, e 68c> Referenced in 69b, 83a.  
 <update f 69a> Referenced in 69b, 83a.  
 <update score for chol 82a> Referenced in 83a.  
 <update score means, lower and upper 82b> Referenced in 83a.  
 <update yp for chol 79c> Referenced in 83a.  
 <update yp for means, lower and upper 80> Referenced in 83a.  
 <upper scores 77d> Referenced in 78a.  
 <vec trick 42a> Referenced in 3.

⟨ W length 71b ⟩ Referenced in 73, 84.

# Bibliography

- Shane Barratt and Stephen Boyd. Covariance prediction via convex optimization. *Optimization and Engineering*, 24(3):2045–2078, 2023. doi: 10.1007/s11081-022-09765-w. [116](#)
- Alan Genz. Numerical computation of multivariate normal probabilities. *Journal of Computational and Graphical Statistics*, 1(2), 1992. doi: 10.1080/10618600.1992.10477010. [1](#), [65](#)
- Alan Genz and Frank Bretz. Methods for the computation of multivariate  $t$  probabilities. *Journal of Computational and Graphical Statistics*, 11(4):950–971, 2002. doi: 10.1198/106186002394. [1](#), [63](#)
- Ivan Matić, Radoš Radoičić, and Dan Stefanica. A sharp Pólya-based approximation to the normal CDF. *Applied Mathematics and Computation*, 322:111–122, 2018. doi: 10.2139/ssrn.2842681. [69](#)