

Elementary volesti tutorial

Vissarion Fisikopoulos

19 April 2024

Abstract

This tutorial has been designed for the Computational Geometry course of the department of Informatics and Telecommunications at National and Kapodistrian University of Athens.

Contents

R language basics	1
The volesti package	4
Sampling	4
Sampling via random walks	6
Volume computation	13
Volume of Birkhoff polytopes	15
Rounding	15
Integration	16
Counting linear extensions	16

R language basics

We start with some R warm-up.

```
# Numerical value
```

```
v <- 23.5
```

```
v
```

```
## [1] 23.5
```

```
# Create a vector
```

```
apple <- c('red','green',"yellow")
```

```
apple
```

```
## [1] "red" "green" "yellow"
```

```
# Create a matrix
```

```
M <- matrix( c('a','a','b','c','b','a'), nrow = 2, ncol = 3, byrow = TRUE)
```

```
M
```

```
##      [,1] [,2] [,3]
```

```
## [1,] "a"  "a"  "b"
```


```
## [2,] "c" "b" "a"
```

```
# Removing a variable  
rm(M)
```

Some basic statistics with R: normal and uniform distributions.

```
# Let's generate a vectors of random numbers from a normal distribution  
x1 <- rnorm(10000, mean=0, sd=1)  
# The breaks argument specifies how many bars are in the histogram  
hist(x1, probability=TRUE, breaks = 100)
```

Histogram of x1



```
# pnorm returns the integral from  $-\infty$  to  $q$  of the pdf of the normal  
# distribution where  $q$  is a Z-score  
# a z-score is the number of standard deviations from the mean a data point is  
pnorm(0)
```

```
## [1] 0.5
```

```
pnorm(2)
```

```
## [1] 0.9772499
```

```
# qnorm function is simply the inverse of the cdf, which you can also  
# think of as the inverse of pnorm  
qnorm(0.5)
```


```
## [1] 0
```

```
qnorm(0.98)
```

```
## [1] 2.053749
```

```
# Sample a random variable from the uniform distribution and plot its histogram
x2 <- runif(1000, min = 0, max = 1)
hist(x2, probability=FALSE, breaks = 100)
```

Histogram of x2



```
# Descriptive statistics
```

```
mean(x1)
```

```
## [1] -0.003786281
```

```
sd(x1)
```

```
## [1] 0.9994998
```

```
var(x1)
```

```
## [1] 0.9989998
```

```
min(x1)
```

```
## [1] -3.925456
```

```
max(x1)
```

```
## [1] 4.053848
```

```
median(x1)
```

```
## [1] -0.008083397
```

```
range(x1)
```

```
## [1] -3.925456 4.053848
```

```
quantile(x1)
```

```
##           0%          25%          50%          75%         100%
```

```
## -3.925456212 -0.677826663 -0.008083397 0.669279949 4.053848411
```

```
summary(x1)
```

```
##      Min.   1st Qu.   Median     Mean   3rd Qu.     Max.
## -3.925456 -0.677827 -0.008083 -0.003786 0.669280 4.053848
```

The volesti package

volesti is a C++ package (with an R interface) for computing estimations of volume of polytopes given by a set of points or linear inequalities or Minkowski sum of segments (zonotopes).

There are various algorithms for volume estimation, sampling and rounding polytopes.

We can download the R package from <https://CRAN.R-project.org/package=volesti> or build the development version from <https://github.com/GeomScale/Rvolesti>

A gentle intro to R interface of volesti (version 1.1.2-2) is <https://journal.r-project.org/archive/2021/RJ-2021-077/RJ-2021-077.pdf>

In this tutorial we are using volesti 1.2.0

```
# first load the volesti library
#install.packages('volesti')
library(volesti)
```

```
## Loading required package: Rcpp
```

```
packageVersion("volesti")
```

```
## [1] '1.2.0'
```

You have access to the documentation of volesti functions like volume computation and sampling.

```
help("volume")
help("sample_points")
```

Full documentation for the CRAN version here: <https://cran.r-project.org/web/packages/volesti/volesti.pdf>

Let's try our first volesti command to estimate the volume of a 3-dimensional cube.

```
P <- gen_cube(4, 'H')
print(volume(P))
```

```
## $log_volume
## [1] 2.838507
##
## $volume
## [1] 17.09023
```

What is the exact volume of P? Did we obtain a good estimation?

Sampling

Sampling uniformly in the square.

```
library(ggplot2)

x1<-runif(1000, min = -1, max = 1)
x2<-runif(1000, min = -1, max = 1)
```


```

g<-ggplot(data.frame( x=x1, y=x2 )) + geom_point( aes(x=x, y=y))

g<-g+annotate("path",
  x=cos(seq(0,2*pi,length.out=100)),
  y=sin(seq(0,2*pi,length.out=100)),color="red")+coord_fixed()

plot(g)

```



Can we estimate the volume of the red ball via sampling? Solution: rejection sampling. The following computation illustrates that this will fail in (not so) high dimensions.

```

for (d in 2:20) {
  num_of_points <- 10000
  count_inside <- 0

  points1 <- matrix(nrow=d, ncol=num_of_points)
  for (i in 1:d) {
    x <- runif(num_of_points, min = -1, max = 1)
    for (j in 1:num_of_points) {
      points1[i,j] <- x[j]
    }
  }

  for (i in 1:num_of_points) {
    if (norm(points1[,i], type="2") < 1) {
      count_inside <- count_inside + 1
    }
  }
}

```

```

}
vol_estimation <- count_inside*2^d/num_of_points
vol_exact <- pi^(d/2)/gamma(d/2+1)

cat(d, vol_estimation, vol_exact, abs(vol_estimation- vol_exact)/
vol_exact, "\n")
}

```

```

## 2 3.1384 3.141593 0.001016253
## 3 4.1464 4.18879 0.01011992
## 4 5.032 4.934802 0.01969639
## 5 5.2192 5.263789 0.008470897
## 6 4.9216 5.167713 0.04762509
## 7 4.2624 4.724766 0.09786008
## 8 3.9936 4.058712 0.01604256
## 9 2.6112 3.298509 0.2083696
## 10 2.6624 2.550164 0.04401127
## 11 1.024 1.884104 0.4565056
## 12 1.6384 1.335263 0.2270244
## 13 0 0.9106288 1
## 14 0 0.5992645 1
## 15 0 0.3814433 1
## 16 0 0.2353306 1
## 17 0 0.1409811 1
## 18 0 0.08214589 1
## 19 0 0.0466216 1
## 20 0 0.02580689 1

```

Sampling via random walks



volesti supports various types of random walks

There are two important parameters cost per step and mixing time that affects the accuracy and performance of the walks. Below we illustrate this by choosing different walk steps for each walk while sampling on the 100-dimensional cube.


```

knitr::opts_chunk$set(echo = TRUE, fig.align = "center")
#run in few secs
library(ggplot2)
library(volesti)
for (step in c(1,20,100,150)){
  for (walk in c("CDHR", "RDHR", "BaW")){
    P <- gen_cube(100, 'H')
    points1 = sample_points(P, n = 1000, random_walk = list("walk" = walk, "walk_length" = step))
    g<-plot(ggplot(data.frame( x=points1[1,], y=points1[2,] )) +
geom_point( aes(x=x, y=y, color=walk)) + coord_fixed(xlim = c(-1,1),
ylim = c(-1,1)) + ggtitle(sprintf("walk length=%s", step)))
  }
}


```




walk length=1




walk length=20




walk length=20




walk length=20




walk length=100




walk length=100




walk length=100




walk length=150




walk length=150



walk length=150



```
P <- gen_cube(100, 'H')
points1 = sample_points(P, n = 1000, random_walk = list("walk" = "aBiW", "walk_length" = 1))
g<-plot(ggplot(data.frame( x=points1[1,], y=points1[2,] )) +
geom_point( aes(x=x, y=y, color=walk)) + coord_fixed(xlim = c(-1,1),
ylim = c(-1,1)))
```



Another way to test the quality of the sample is to use some statistical tests, e.g. the effective sample size (ess)

```
P <- gen_cube(100, 'H')
for (walk in c("CDHR", "RDHR", "BaW", "aBiW")){
  points1 = sample_points(P, n = 1000, random_walk = list("walk" = walk, "walk_length" = 10))
  cat(walk, min(ess(points1)), "\n")
}
```

```
## CDHR 18.57934
## RDHR 3.232083
## BaW 3.13437
## aBiW 647.0717
```

Volume computation

Now let's compute our first example. The volume of the 3-dimensional cube.

We want to compare with an exact volume computation software *geometry* (R wrapper for *qhull*) so we focus on *V-polytopes*.

```

#install.packages("geometry")
library(geometry)

PV <- gen_cube(3, 'V')
str(PV)

## Formal class 'Vpolytope' [package "volesti"] with 3 slots
## ..@ V      : num [1:8, 1:3] -1 1 -1 1 -1 1 -1 1 -1 ...
## ..@ volume: num 8
## ..@ type  : chr "Vpolytope"

geom_values <- convhulln(PV@V, options = 'FA')
vol_approx <- volume(PV)

cat(sprintf("exact vol = %f\napprx vol = %f\nrelative error = %f\n",
           geom_values$vol, vol_approx$volume, abs(geom_values$vol-vol_approx$volume)/geom_values$vol))

## exact vol = 8.000000
## aprrx vol = 7.781443
## relative error = 0.027320

```

Now try a higher dimensional example. By setting the error parameter we can control the approximation of the algorithm.

```

PH = gen_cube(10, 'H')
volumes <- list()
for (i in 1:10) {
  volumes[[i]] <- volume(PH, settings = list("error" = 1))$volume # default parameters
}
options(digits=10)
summary(as.numeric(volumes))

```

```

##      Min.   1st Qu.   Median     Mean   3rd Qu.     Max.
## 953.1138 1020.5465 1038.6450 1050.1363 1081.8522 1137.6432

```

```

volumes <- list()
for (i in 1:10) {
  volumes[[i]] <- volume(PH, settings = list("error" = 0.01))$volume
}
summary(as.numeric(volumes))

```

```

##      Min.   1st Qu.   Median     Mean   3rd Qu.     Max.
## 969.6038 994.4964 1011.0135 1011.1967 1028.7704 1067.2527

```

Deterministic algorithms for volume are limited to low dimensions (e.g. less than 15)

```

library(geometry)

P = gen_rand_vpoly(15, 20)
# this will return an error about memory allocation, i.e. the dimension is too high for ghull
#tim1 <- system.time({ geom_values = convhulln(P$V, options = 'FA') })

time <- system.time({ vol <- volume(P)$volume })
print(vol)

## [1] 9.835096031e-13

```

```
print(time)
```

```
## user system elapsed
## 3.390 0.345 3.775
```

Volume of Birkhoff polytopes

We now continue with a more interesting example, the 10-th Birkhoff polytope. It is known from <https://arxiv.org/pdf/math/0305332.pdf> that its volume equals

$$\text{vol}(B_{10}) = \frac{727291284016786420977508457990121862548823260052557333386607889}{828160860106766855125676318796872729344622463533089422677980721388055739956270293750883504892820848640000000}$$
 obtained via massive parallel computation.

```
library(volesti)
```

```
exact <- 727291284016786420977508457990121862548823260052557333386607889/828160860106766855125676318796872729344622463533089422677980721388055739956270293750883504892820848640000000
exact
```

```
## [1] 8.782005031e-46
```

```
# warning the following will take around a minute
```

```
B <- gen_birkhoff(10)
time <- system.time({ vol <- volume(B)$volume })
print(vol)
```

```
## [1] 7.680032381e-55
```

```
print(time)
```

```
## user system elapsed
## 35.006 0.121 35.126
```

Compare our computed estimation with the “normalized” floating point version of $\text{vol}(B_{10})$

```
n <- 10
vol_B10 <- 727291284016786420977508457990121862548823260052557333386607889/828160860106766855125676318796872729344622463533089422677980721388055739956270293750883504892820848640000000
print(vol_B10/(n^(n-1)))
```

```
## [1] 8.782005031e-55
```

Rounding

We generate skinny polytopes, in particular skinny cubes of the form $\{x = x_1, \dots, x_d \mid x_1 \leq 100, x_1 \geq -100, x_i \leq 1, x_i \geq -1, x_i \in \mathbb{R}, \text{ for } i = 2, \dots, d\}$ Random walks perform poorly on those polytopes especially as the dimension increases.

```
library(ggplot2)
```

```
d <- 100
```

```
P = gen_skinny_cube(d)
P <- rotate_polytope(P)$P
```

```
for (walk in c("CDHR", "RDHR", "BaW")){
  points1 = sample_points(P, n = 1000, random_walk = list("walk" = walk, "walk_length" = 100))
  cat(walk, min(ess(points1)), "\n")
}
```

```
## CDHR 9.027373
## RDHR 3.426704
## BaW 8.221127

points1 = sample_points(P, n = 1000, random_walk = list("walk" = "aBiW", "walk_length" = 1))
cat(walk, min(ess(points1)), "\n")
```

```
## BaW 12.13126
```

Applying a rounding algorithm improves the convergence of walks.

```
Pr <- round_polytope(P)$P
for (walk in c("CDHR", "RDHR", "BaW")){
  points1 = sample_points(Pr, n = 1000, random_walk = list("walk" = walk, "walk_length" = 100))
  cat(walk, min(ess(points1)), "\n")
}
```

```
## CDHR 13.03921
## RDHR 15.38379
## BaW 4.842095
```

```
points1 = sample_points(Pr, n = 1000, random_walk = list("walk" = "aBiW", "walk_length" = 1))
cat(walk, min(ess(points1)), "\n")
```

```
## BaW 284.9732
```

Integration

We can use sampling and volume estimation to estimate integrals over polyhedral domains. Below there is an example with a degree 2 polynomial over a 3-dimensional cube.

```
library(cubature) # load the package "cubature"
f <- function(x) { 2/3 * (2 * x[1]^2 + x[2] + x[3]) + 10 } # "x" is vector
adaptIntegrate(f, lowerLimit = c(-1, -1, -1), upperLimit = c(1, 1, 1))$integral
```

```
## [1] 83.55556
```

```
# Simple Monte Carlo integration
# https://en.wikipedia.org/wiki/Monte_Carlo_integration
P = gen_cube(3, 'H')
num_of_points <- 10000
points1 <- sample_points(P, random_walk = list("Walk" = "aBiW", "walk_length" = 100), n=num_of_points)
int<-0
for (i in 1:num_of_points){
  int <- int + f(points1[,i])
}
V <- volume(P)$volume
print(int*V/num_of_points)
```

```
## [1] 85.63327
```

Counting linear extensions

Let $G = (V, E)$ be an acyclic digraph with $V = [n] := \{1, 2, \dots, n\}$. One might want to consider G as a representation of the partially ordered set (poset) $V : i \geq j$ if and only if there is a directed path from node i to node j . A permutation π of $[n]$ is called a linear extension of G (or the associated poset V) if $\pi^{-1}(i) \geq \pi^{-1}(j)$ for every edge $i \rightarrow j \in E$.

