

Statistics 222, Spatial Statistics.

Outline for the day:

1. Variance of compound Poisson and simulation.
2. Hawkes process.
3. ETAS.
4. Likelihood.
5. Maximum likelihood estimation.
6. R.

1. Variance of the compound Poisson processes, from last time.

Fix B. Let M denote M(B). For a compound Poisson process,

$$\begin{aligned}
 E N(B) &= \sum E(N(B)|m) f(m), \text{ where the sum is from } m = 0, 1, 2, \dots, \\
 &= \sum E(Z_1 + Z_2 + \dots + Z_m) f(m) \\
 &= \sum (m E(Z)) f(m) \\
 &= E(Z) \sum m f(m) \\
 &= E(Z) E(M) = c|B| E(Z).
 \end{aligned}$$

$$\begin{aligned}
 E(N(B)^2) &= \sum E(N(B)^2|m) f(m) \\
 &= \sum E(Z_1 + Z_2 + \dots + Z_m)^2 f(m) \\
 &= \sum (mE(Z^2) + (m^2-m) E(Z)^2)f(m) \\
 &= E(Z^2)\sum mf(m) - E(Z)^2 \sum m f(m) + E(Z)^2 \sum m^2 f(m) \\
 &= E(Z^2) E(M) - E(Z)^2 E(M) + E(Z)^2 E(M^2) \\
 &= V(Z)E(M) + E(Z)^2 E(M^2).
 \end{aligned}$$

$$\begin{aligned}
 \text{So } V(N(B)) &= E(N(B)^2) - (E(N(B)))^2 \\
 &= V(Z)E(M) + E(Z)^2 E(M^2) - E(M)^2 E(Z)^2 \\
 &= V(Z) E(M) + E(Z)^2 (E(M^2) - E(M)^2) \\
 &= V(Z) E(M) + E(Z)^2 V(M).
 \end{aligned}$$

M is Poisson, so $E(M) = V(M) = c|B|$, so

$$V(N(B)) = c|B| (V(Z) + E(Z)^2) = c|B| E(Z^2) \geq E N(B), \text{ since } E(Z^2) \geq E(Z).$$

Simulation.

One can simulate spatial-temporal point processes by *thinning*.

Lewis, P. and Shedler, G. (1979). Simulation of nonhomogeneous poisson processes by thinning. *Naval Research Logistics Quarterly*, 26:403–413, 1979.



Jesper Møller

Suppose λ has some upper bound, B . $\lambda(t, x, y) \leq B$ everywhere.

First, simulate a stationary Poisson process N with intensity B .

For $i = 1, 2, \dots$ keep point τ_i with probability $\lambda(\tau_i)/B$. We saw this in the code for Day 3 for simulating inhomogeneous Poisson processes and it works for other processes too.

Boundary issues can be important in simulation. For Gibbs processes, for instance, the simulation can be biased because of missing points outside the observation region. For Hawkes processes, the simulation will tend to be biased by having too few points at the beginning of the simulation. One can have burn-in, by simulating points outside the observation region or before time 0, or in some cases some fancy weighting schemes can be done to achieve *perfect* simulation without burn-in.

Møller, J. and Waagepetersen, R. (2003). *Statistical Inference and Simulation for Spatial Point Processes*. Chapman and Hall, Boca Raton.


2. Hawkes process.

A Hawkes process or *self-exciting* process has conditional intensity

$$\begin{aligned}\lambda(t,x,y) &= \mu(x,y) + \kappa \int_{t' < t} g(t-t', x-x', y-y') dN(t', x', y') \\ &= \mu(x,y) + \kappa \sum_{\{t', x', y': t' < t\}} g(t-t', x-x', y-y').\end{aligned}$$

g is called the *triggering function* or *triggering density* and κ is the *productivity*.

If g is a density function, then κ is the expected number of points triggered directly by each point. Each background point, associated with $\mu(x,y)$, is expected to generate $\kappa + \kappa^2 + \kappa^3 + \dots = 1/(1-\kappa) - 1$ triggered points, so the exp. fraction of background pts is $1-\kappa$.



3. ETAS process.

An *Epidemic-Type Aftershock Sequence (ETAS)* process is a marked version of the Hawkes process, where points have different productivities depending on their magnitudes. Ogata (1988, 1998) introduced

$$\lambda(t,x,y)=\mu(x,y) + \sum_{\{t',x',y': t' < t\}} g(t-t',x-x',y-y')h(m'),$$

where $\mu(x,y)$ is estimated by smoothing observed large earthquakes,

$$h(m) = \kappa e^{\alpha(m-m_0)},$$

where m_0 is the catalog cutoff magnitude, and $g(t,x,y) = g_1(t) g_2(r^2)$, where $r^2 = \|(x,y)\|^2$, and g_1 and g_2 are power-law or *Pareto* densities, $g_1(t) = (p-1) c^{p-1} (t+c)^{-p}$. $g_2(r^2) = (q-1) d^{q-1} (r^2+d)^{-q}$.

An alternative is where g_2 is exponential or sum of exponentials.



Yoshihiko Ogata




Figure 3: Recorded epicenters of Hector Mine $M \geq 3.0$ earthquakes from 10/16/1999 to 12/23/2000, from SCSN.

4. Likelihood.

For iid draws t_1, t_2, \dots, t_n , from some density $f(\theta)$, the likelihood is simply

$$L(\theta) = f(t_1; \theta) \times f(t_2; \theta) \times \dots \times f(t_n; \theta) \\ = \prod f(t_i; \theta).$$

This is the probability density of observing $\{t_1, t_2, \dots, t_n\}$, as a function of the parameter θ .



For a nonstationary Poisson process with intensity $\lambda(\theta)$, on $[0, T]$, the likelihood of observing the points $\{\tau_1, \tau_2, \dots, \tau_n\}$ is simply

$$\lambda(\tau_1) \times \lambda(\tau_2) \times \dots \times \lambda(\tau_n) \times \\ \exp\{-A(\tau_1)\} \times \exp\{-(A(\tau_2)-A(\tau_1))\} \times \dots \times \exp\{-(A(T)-A(\tau_n))\}, \\ = \prod \lambda(\tau_i) \exp\{-A(T)\},$$

where $A(t) = \int_0^t \lambda(t) dt$.

$P\{k \text{ points in } (\tau_1, \tau_2)\}$ is $\exp(-B) B^k/k! = \exp(-B)$ for $k = 0$, where $B = \int_{\tau_1}^{\tau_2} \lambda(t) dt$.


Likelihood, continued.

For a nonstationary Poisson process with intensity $\lambda(t)$,
on $[0, T]$, the likelihood of observing the points $\{\tau_1, \tau_2, \dots, \tau_n\}$ is simply
 $\lambda(\tau_1) \times \lambda(\tau_2) \times \dots \times \lambda(\tau_n) \times$

$$\exp\{-A(\tau_1)\} \times \exp\{-(A(\tau_2)-A(\tau_1))\} \times \dots \times \exp\{-(A(T)-A(\tau_n))\},$$
$$= \prod \lambda(\tau_i) \exp\{-A(T)\},$$

where $A(t) = \int_0^t \lambda(t) dt$.

$P\{k \text{ points in } (\tau_1, \tau_2)\}$ is $\exp(-B) B^k/k! = \exp(-B)$ for $k = 0$,
where $B = \int_{\tau_1}^{\tau_2} \lambda(t) dt$.



So the log likelihood is $\sum \log(\lambda(\tau_i)) - A(T)$.

In the spatial-temporal case, the log likelihood is simply
 $\sum \log(\lambda(\tau_i)) - \int \lambda(t, x, y) dt dx dy$.

5. Maximum likelihood estimation.

Find $\hat{\theta}$ ($= \theta^*$) maximizing $l(\theta) = \sum \log(\lambda(\tau_i)) - \int \lambda(t, x, y) dt dx dy$.

Ogata (1978) showed that the resulting estimate, $\hat{\theta}$, is, under standard conditions, asymptotically unbiased, $E(\hat{\theta}) \rightarrow \theta$, consistent, $P(|\hat{\theta} - \theta| > \varepsilon) \rightarrow 0$ as $T \rightarrow \infty$, for any $\varepsilon > 0$, asymptotically normal, $\hat{\theta} \rightarrow_D \text{Normal}$ as $T \rightarrow \infty$, and asymptotically efficient, min. variance among asymptotically unbiased estimators.

Further, he showed standard errors for $\hat{\theta}$ can be constructed using the diagonal elements of the inverse of the Hessian of L evaluated at $\hat{\theta}$.
`sqrt(diag(solve(loglikelihood$hess)))`



Ogata, Y. (1978). The asymptotic behaviour of maximum likelihood estimators for stationary point processes. Ann. Inst. Statist. Math. 30, 243-261.

Maximum likelihood estimation continued.

The conditions of Ogata (1978) can be relaxed a bit for Poisson processes [1], and for certain spatial-temporal process in general [2].

Even if the process is not Poisson, under some circumstances [3] the parameters governing the unconditional intensity, $E\lambda$, can be consistently estimated by maximizing $L_P(\theta) = \sum \log(E\lambda(\tau_i)) - \int E\lambda(t,x,y) dt dx dy$. Basically pretend the process is Poisson.

Suppose you are missing some covariate that might affect λ . Under general conditions, the MLE will nevertheless be consistent, provided the effect of the missing covariate is small [4].

[1] Rathbun, S.L., and Cressie, N. (1994). Asymptotic properties of estimators for the parameters of spatial inhomogeneous Poisson point processes. *Adv. Appl. Probab.* 26, 122–154.

[2] Rathbun, S.L., (1996). Asymptotic properties of the maximum likelihood estimator for spatio-temporal point processes. *JSPI* 51, 55–74.

[3] Schoenberg, F.P. (2004). Consistent parametric estimation of the intensity of a spatial-temporal point process. *JSPI* 128(1), 79--93.

[4] Schoenberg, F.P. (2016). A note on the consistent estimation of spatial-temporal point process parameters. *Statistica Sinica*, 26, 861-879.

Maximum likelihood estimation continued.

λ is completely separable if $\lambda(t, x, y; \theta) = \theta_3 \lambda_0(t; \theta_0) \lambda_1(t, x; \theta_1) \lambda_2(t, y; \theta_2)$.

Suppose N has marks too. λ is separable in mark (or coordinate) i if

$$\lambda(t, x, y, m_1, m_2, \dots, m_k; \theta) = \theta_2 \lambda_i(t, m_i; \theta_i) \lambda_{-i}(t, x, y, m_{-i}; \theta_{-i}).$$

Suppose you are neglecting some *mark* or coordinate of the process. Under some conditions, the MLE of the other parameters will nevertheless be consistent [1].

In maximizing $L(\theta) = \sum \log(\lambda(\tau_i)) - \int \lambda(t, x, y) dt dx dy$, it is typically straightforward to compute the sum, but the integral can be tricky esp. when the conditional intensity is very volatile. One trick noted in [2] is that, for a Hawkes process where $\lambda(t, x, y) = \mu(x, y) + \kappa \sum_{\{t', x', y': t' < t\}} g(t-t', x-x', y-y')$, where g is a density, and $\int \mu(x, y) dx dy = \mu$,

$$\begin{aligned} \int \lambda(t, x, y) dt dx dy &= \mu T + \kappa \int \sum g(t-t', x-x', y-y') dt dx dy \\ &= \mu T + \kappa \sum \int g(t-t', x-x', y-y') dt dx dy \\ &\sim \mu T + \kappa N. \end{aligned}$$

[1] Schoenberg, F.P. (2016). A note on the consistent estimation of spatial-temporal point process parameters. *Statistica Sinica*, 26, 861-879.

[2] Schoenberg, F.P. (2013). Facilitated estimation of ETAS. *Bulletin of the Seismological Society of America*, 103(1), 601-605.

6. Questions.

The difference between ETAS and a Hawkes process is ...

- a) an ETAS process is more strongly clustered.
- b) the points of an ETAS process occur at different locations.
- c) the points of an ETAS process have different productivities.
- d) the points of an ETAS process have different triggering functions.

Questions.

The difference between ETAS and a Hawkes process is ...

- a) an ETAS process is more strongly clustered.
- b) the points of an ETAS process occur at different locations.
- c) the points of an ETAS process have different productivities.**
- d) the points of an ETAS process have different triggering functions.

Which of the following can possibly have two points within distance .01 of each other?

- a) a hardcore process with $\sigma = .01$.
- b) a Strauss process with $R = .01$.
- c) a Matern I process with $r = .01$.
- d) a Matern II process with $r = .01$.

Questions.

The difference between ETAS and a Hawkes process is ...

- a) an ETAS process is more strongly clustered.
- b) the points of an ETAS process all occur at different locations.
- c) the points of an ETAS process all have different productivity.**
- d) the points of an ETAS process all have different triggering functions.

Which of the following can possibly have two points within distance .01 of each other?

- a) a hardcore process with $\sigma = .01$.
- b) a Strauss process with $R = .01$.**
- c) a Matern I process with $r = .01$.
- d) a Matern II process with $r = .01$.

Code.

Cox process driven by a Brownian Sheet.

```
m = 100
m1 = m-1
y = matrix(rep(0,m*m),ncol=m)
for(i in c(2:m)) y[1,i] = y[1,(i-1)] + rnorm(1)/2
for(i in c(2:m)) y[i,1] = y[(i-1),1] + rnorm(1)/2
for(i in c(2:m)){
  for(j in c(2:m)){
    y[i,j] = (y[(i-1),j] + rnorm(1))/2 + (y[i,(j-1)] + rnorm(1))/2
  }
}
x1 = c(0:m1)
x2 = c(0:m1)
x3 = y[(2:m),(2:m)]-y[1,1]; x3 = (x3-min(x3))/(max(x3)-min(x3))/(m^2)*100
x4 = (0:100)/100*(max(x3)-min(x3))+min(x3)
par(mfrow=c(1,2))
image(x1,x2,x3,col=gray((64:0)/64),xlab="column",ylab="row")
for(i in 1:m1){for(j in 1:m1){
  n = rpois(1,x3[i,j]); if(n>0){x = runif(n)+i-1;y=runif(n)+j-1;points(x,y,col="red")}} }
plot(c(0,10),c(0.8*min(x4),1.2*max(x4)),type="n",axes=F,xlab="",ylab="")
image(c(-1:1),x4,matrix(rep(x4,2),ncol=101,byrow=T),add=T,col=gray((64:0)/64))
text(2,min(x4),as.character(signif(min(x4),2)),cex=1)
text(2,(max(x4)+min(x4))/2,as.character(signif((max(x4)+min(x4))/2,2)),cex=1)
text(2,max(x4),as.character(signif(max(x4),2)),cex=1)
```

Code.

```
install.packages("spatstat")  
library(spatstat)
```

```
## STRAUSS process  
z = rStrauss(100,0.7,0.05)  
plot(z, pch=2,cex=.5)
```

```
## HARDCORE process  
z = rHardcore(100,0.05)  
plot(z, pch=2,cex=.5)
```

```
## MATERN(I).  
z = rMaternI(20,.05)  
plot(z, pch=2,cex=.5)
```


Code.

```
## MATERN(II)
z = rMaternII(100,.05)
plot(z,pch=2,cex=.5)

## HAWKES.
install.packages("hawkes")
library(hawkes)


lambda0 = c(0.2,0.2)

alpha = matrix(c(0.5,0,0,0.5),byrow=TRUE,nrow=2)
beta = c(0.7,0.7)

horizon = 3600

h = simulateHawkes(lambda0,alpha,beta,horizon)
plot(c(0,3600),c(0,3),type="n",xlab="t",ylab="x")
points(h[[1]],.5+runif(length(h[[1]])),pch=2,cex=.1)
points(h[[2]],1.5+runif(length(h[[2]])),pch=3,cex=.1)
```

R. To download and install R, start at www.r-project.org, click on “download R”, scroll down and click on one of the mirrors closest to you. From there, click on “download R for ...”, and then get the latest version.



[Home]

Download

[CRAN](#)

R Project

[About R](#)

[Logo](#)

[Contributors](#)

[What's New?](#)

[Reporting Bugs](#)

[Conferences](#)

[Search](#)

[Get Involved: Mailing](#)

[Lists](#)

[Developer Pages](#)

[R Blog](#)

R Foundation

The R Project for Statistical Computing

Getting Started


R is a free software environment for statistical computing and graphics. It compiles and runs on a wide variety of UNIX platforms, Windows and MacOS. To **download R**, please choose your preferred [CRAN mirror](#).

If you have questions about R like how to download and install the software, or what the license terms are, please read our [answers to frequently asked questions](#) before you send an email.

News

- **R version 4.0.5 (Shake and Throw)** has been released on 2021-03-31.
- Thanks to the organisers of useR! 2020 for a successful online conference. Recorded tutorials and talks from the conference are available on the [R Consortium YouTube channel](#).
- **R version 3.6.3 (Holding the Windsock)** was released on 2020-02-29.
- You can support the R Foundation with a renewable subscription as a [supporting member](#)

R. To download and install *R*, start at www.r-project.org, click on “download *R*”, scroll down and click on one of the mirrors closest to you. From there, click on “download R for ...”, and then get the latest version.



CRAN Mirrors

The Comprehensive R Archive Network is available at the following URLs, please choose a location close to you. Some statistics on the status of the mirrors can be found here: [main page](#), [windows release](#), [windows old release](#).

If you want to host a new mirror at your institution, please have a look at the [CRAN Mirror HOWTO](#).

0-Cloud

<https://cloud.r-project.org/>

Automatic redirection to servers worldwide, currently sponsored by Rstudio

Algeria

<https://cran.usthb.dz/>

University of Science and Technology Houari Boumediene

Argentina

<http://mirror.fcaglp.unlp.edu.ar/CRAN/>

Universidad Nacional de La Plata

Australia

<https://cran.csiro.au/>

CSIRO

<https://mirror.aarnet.edu.au/pub/CRAN/>

AARNET

<https://cran.ms.unimelb.edu.au/>

School of Mathematics and Statistics, University of Melbourne

<https://cran.curtin.edu.au/>

Curtin University

Austria

<https://cran.wu.ac.at/>

Wirtschaftsuniversität Wien

Belgium

<https://www.freeststatistics.org/cran/>

Patrick Wessa

<https://ftp.belnet.be/mirror/CRAN/>

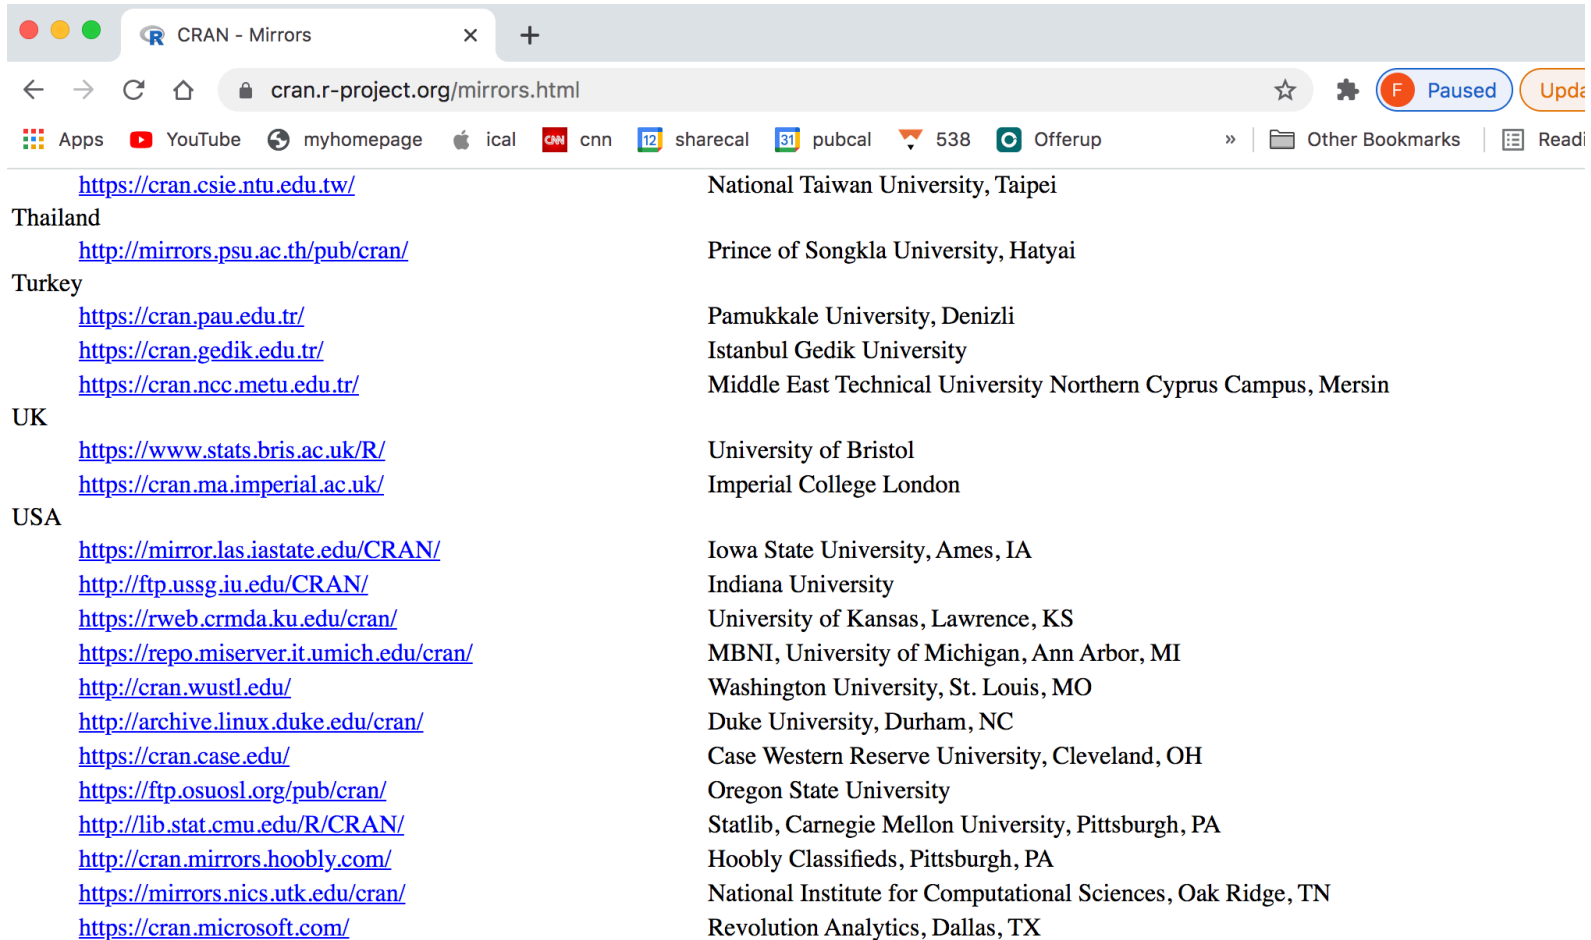
Belnet, the Belgian research and education network

Brazil

<https://nbcgib.uesc.br/mirrors/cran/>

Computational Biology Center at Universidade Estadual de Santa Cruz

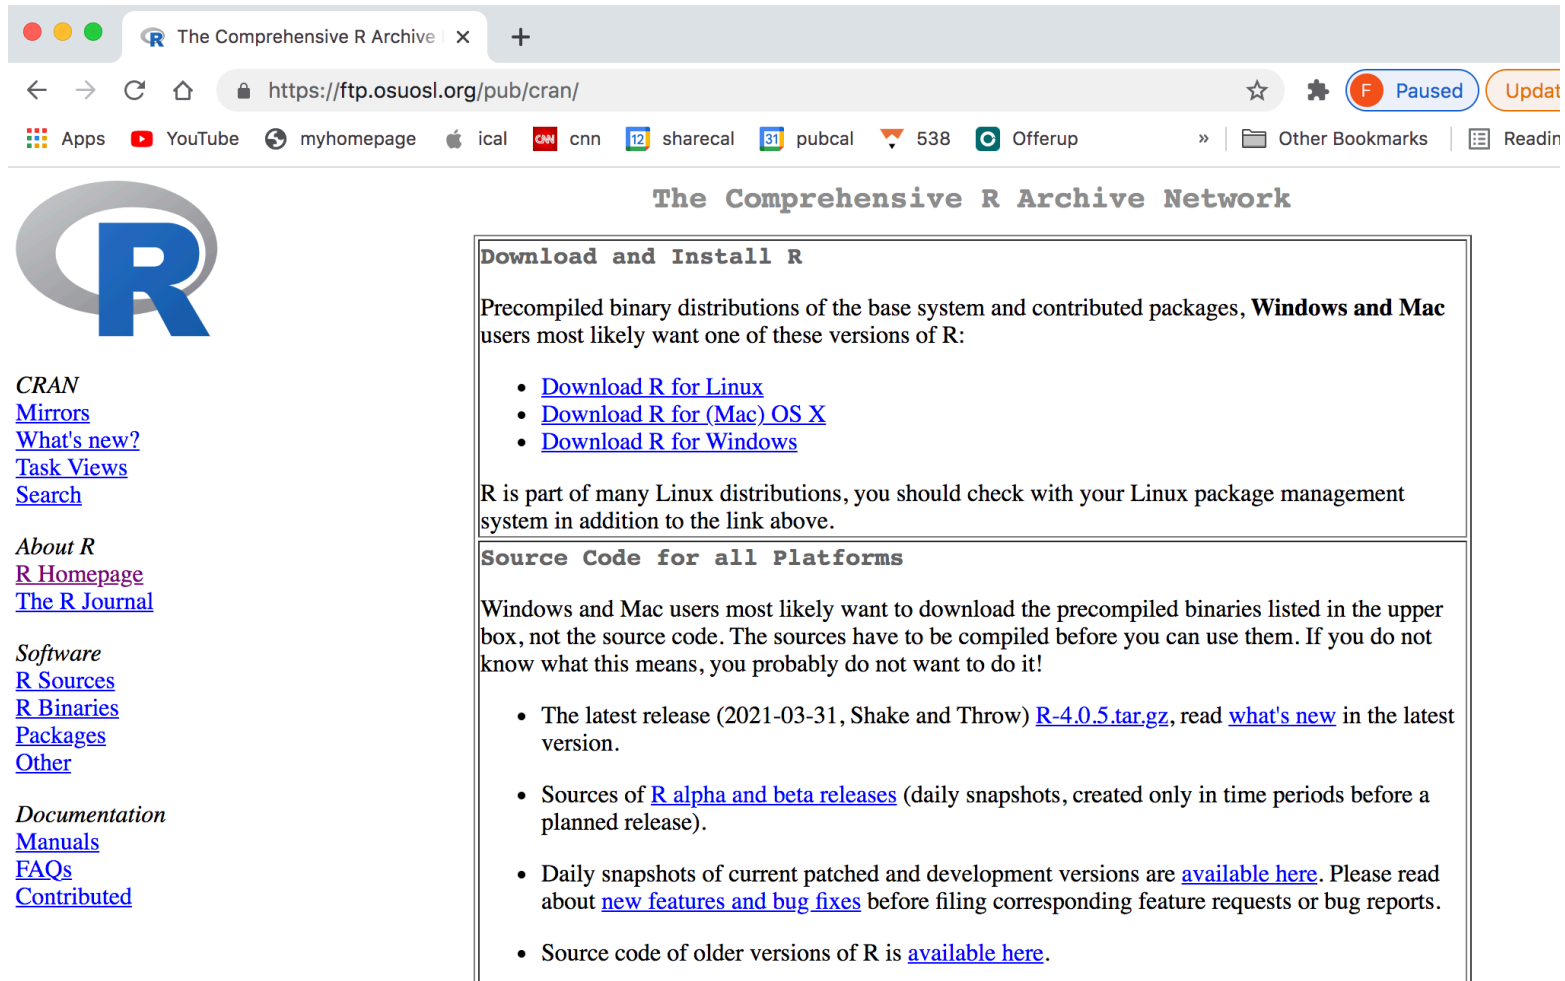
R. To download and install R, start at www.r-project.org, click on “download R”, scroll down and click on one of the mirrors closest to you. From there, click on “download R for ...”, and then get the latest version.



The screenshot shows a web browser window with the address bar displaying cran.r-project.org/mirrors.html. The page content lists mirrors for various countries, including Thailand, Turkey, UK, and USA. Each country has a list of URLs and the names of the hosting institutions.

Country	URL	Institution
Thailand	https://cran.csie.ntu.edu.tw/	National Taiwan University, Taipei
	http://mirrors.psu.ac.th/pub/cran/	Prince of Songkla University, Hatyai
Turkey	https://cran.pau.edu.tr/	Pamukkale University, Denizli
	https://cran.gedik.edu.tr/	Istanbul Gedik University
	https://cran.ncc.metu.edu.tr/	Middle East Technical University Northern Cyprus Campus, Mersin
UK	https://www.stats.bris.ac.uk/R/	University of Bristol
	https://cran.ma.imperial.ac.uk/	Imperial College London
USA	https://mirror.las.iastate.edu/CRAN/	Iowa State University, Ames, IA
	http://ftp.ussg.iu.edu/CRAN/	Indiana University
	https://rweb.crmda.ku.edu/cran/	University of Kansas, Lawrence, KS
	https://repo.miserver.it.umich.edu/cran/	MBNI, University of Michigan, Ann Arbor, MI
	http://cran.wustl.edu/	Washington University, St. Louis, MO
	http://archive.linux.duke.edu/cran/	Duke University, Durham, NC
	https://cran.case.edu/	Case Western Reserve University, Cleveland, OH
	https://ftp.osuosl.org/pub/cran/	Oregon State University
	http://lib.stat.cmu.edu/R/CRAN/	Statlib, Carnegie Mellon University, Pittsburgh, PA
	http://cran.mirrors.hoobly.com/	Hoobly Classifieds, Pittsburgh, PA
	https://mirrors.nics.utk.edu/cran/	National Institute for Computational Sciences, Oak Ridge, TN
	https://cran.microsoft.com/	Revolution Analytics, Dallas, TX

R. To download and install R, start at www.r-project.org, click on “download R”, scroll down and click on one of the mirrors closest to you. From there, click on “download R for ...”, and then get the latest version.



The Comprehensive R Archive Network

Download and Install R

Precompiled binary distributions of the base system and contributed packages, **Windows and Mac** users most likely want one of these versions of R:

- [Download R for Linux](#)
- [Download R for \(Mac\) OS X](#)
- [Download R for Windows](#)

R is part of many Linux distributions, you should check with your Linux package management system in addition to the link above.

Source Code for all Platforms

Windows and Mac users most likely want to download the precompiled binaries listed in the upper box, not the source code. The sources have to be compiled before you can use them. If you do not know what this means, you probably do not want to do it!

- The latest release (2021-03-31, Shake and Throw) [R-4.0.5.tar.gz](#), read [what's new](#) in the latest version.
- Sources of [R alpha and beta releases](#) (daily snapshots, created only in time periods before a planned release).
- Daily snapshots of current patched and development versions are [available here](#). Please read about [new features and bug fixes](#) before filing corresponding feature requests or bug reports.
- Source code of older versions of R is [available here](#).

CRAN
[Mirrors](#)
[What's new?](#)
[Task Views](#)
[Search](#)


About R
[R Homepage](#)
[The R Journal](#)

Software
[R Sources](#)
[R Binaries](#)
[Packages](#)
[Other](#)

Documentation
[Manuals](#)
[FAQs](#)
[Contributed](#)

R. To download and install R, start at www.r-project.org, click on “download R”, scroll down and click on one of the mirrors closest to you.

From there, click on “download R for ...”, and then get the latest version.



[CRAN](#)
[Mirrors](#)
[What's new?](#)
[Task Views](#)
[Search](#)

[About R](#)
[R Homepage](#)
[The R Journal](#)

[Software](#)
[R Sources](#)
[R Binaries](#)
[Packages](#)
[Other](#)

[Documentation](#)
[Manuals](#)
[FAQs](#)
[Contributed](#)

R for Mac OS X

This directory contains binaries for a base distribution and packages to run on Mac OS X (release 10.6 and above). Mac OS 8. to 9.2 (and Mac OS X 10.1) are no longer supported but you can find the last supported release of R for these systems (which is R 1.7.1) [here](#). Releases for old Mac OS X systems (through Mac OS X 10.5) and PowerPC Macs can be found in the [old](#) directory.

Note: CRAN does not have Mac OS X systems and cannot check these binaries for viruses. Although we take precautions when assembling binaries, please use the normal precautions with downloaded executables.

Package binaries for R versions older than 3.2.0 are only available from the [CRAN archive](#) so users of such versions should adjust the CRAN mirror setting (<https://cran-archive.r-project.org>) accordingly.

R 4.0.5 "Shake and Throw" released on 2021/03/31

Please check the SHA1 checksum of the downloaded image to ensure that it has not been tampered with or corrupted during the mirroring process. For example type

```
openssl sha1 R-4.0.5.pkg
```

in the *Terminal* application to print the SHA1 checksum for the R-4.0.5.pkg image. On Mac OS X 10.7 and later you can also validate the signature using

```
pkgutil --check-signature R-4.0.5.pkg
```

Latest release:

[R-4.0.5.pkg](#) (notarized and signed)
SHA1-
hash: 2f683b3c10f1a9aad92723663abef02285b6132
(ca. 85MB)

R 4.0.5 binary for macOS 10.13 (High Sierra) and higher, signed and notarized package. Contains R 4.0.5 framework, R.app GUI 1.74 in 64-bit for Intel Macs, Tcl/Tk 8.6.6 X11 libraries and Texinfo 6.7. The latter two components are optional and can be omitted when choosing "custom install", they are only needed if you want to use the `tcltk` R package or build package documentation from sources.