

A software review for extreme value analysis

Eric Gilleland · Mathieu Ribatet ·
Alec G. Stephenson

Received: 8 September 2011 / Revised: 12 June 2012 /
Accepted: 27 June 2012 / Published online: 20 July 2012
© The Author(s) 2012. This article is published with open access at Springerlink.com

Abstract Extreme value methodology is being increasingly used by practitioners from a wide range of fields. The importance of accurately modeling extreme events has intensified, particularly in environmental science where such events can be seen as a barometer for climate change. These analyses require tools that must be simple to use, but must also implement complex statistical models and produce resulting inferences. This document presents a review of the software that is currently available to scientists for the statistical modeling of extreme events. We discuss all software known to the authors, both proprietary and open source, targeting different data types and application areas. It is our intention that this article will simplify the process of understanding the available software, and will help promote the methodology to an expansive set of scientific disciplines.

Keywords Extreme value theory · Software development · Spatial extremes · Statistical computing

1 Introduction

In the previous five years the development of software for statistical extremes has been rapid, particularly in the open source environment, where individual academics

E. Gilleland (✉)
Research Applications Laboratory, National Center for Atmospheric Research, Boulder, CO, USA
e-mail: ericg@ucar.edu

M. Ribatet
Institute of Mathematics, University of Montpellier II, Montpellier, France

A. G. Stephenson
CSIRO, Mathematics, Informatics and Statistics, Clayton South, Melbourne, VIC, Australia

have made available more tools than ever before. The motivation derives from both theoretical research and practical applications, particularly in environmental science. Perhaps there is also increasing recognition that for theoretical research to be applied in practical applications, it must be easily reproducible, and the creation of software is a primary means of achieving this.

The pace of progress has a cost, namely that there is now more software available than ever, and consequently some effort must be made in finding the correct tool for a particular job. The amount will only increase in future years, so the challenge is to review their coverage and categorize them appropriately in order to ease the load for the end-user. This article can be seen as a part of the solution to the potential burden of choice. There are several statistical techniques that have more than one implementation; we emphasize that our intention is not a critical comparison, but a summary of available utility.

In Section 2 we briefly give the background to some models that form a basis for the theory of statistical extremes, in order to put the following summaries of software packages in context. Section 3 describes the software, beginning with packages available in the open source statistical environment R (R Development Core Team 2012), and ending with stand-alone software and software designed for other environments such as MatLab and S-Plus. A concluding discussion is given in Section 4.

2 Background and notation

It is assumed here that the reader is familiar with statistical extreme value analysis, and those who are not are referred to texts on the subject (e.g., Beirlant et al. 2004; Coles 2001; de Haan and Ferreira 2006; Embrechts et al. 1997; Reiss and Thomas 2001). In this section, we give a very brief background primarily to identify the notation used here.

The generalized extreme value (GEV) family of distribution functions has theoretical support for fitting to block maximum data whereby the blocks are sufficiently large, and is given by

$$G(z) = \exp \left\{ - \left[1 + \xi \frac{z - \mu}{\sigma} \right]_+^{-1/\xi} \right\}, \quad (2.1)$$

where $\mu \in \mathbb{R}$, $\sigma > 0$ and $\xi \in \mathbb{R}$ are location, scale and shape parameters, respectively, and $y_+ = \max\{y, 0\}$. There are assumptions implicit in fitting the GEV distribution to data, and such assumptions should be checked (e.g., by examining the appropriate qq-plot, checking the data for dependence, etc.).

Similar theory holds for excesses over a high threshold u , whereby the generalized Pareto (GP) distribution family is now supported. The GP distribution has two parameters (scale and shape), and an approximate equivalence between the tails of the GEV and GP distributions exists. We again refer to these parameters as σ and ξ , respectively, noting that $\sigma > 0$ and $\xi \in \mathbb{R}$. Collectively, we refer to these distributions as extreme value distributions (EVD). A point process characterization is

possible in this context, whereby the frequencies of exceeding u are also modeled, and the excesses themselves follow an EVD. Specifically, such processes can be modeled as non-homogeneous Poisson processes. One characterization is a special case of a marked Poisson process with a mark (or excess) being associated with each exceedance, and modeled by the GP distribution. Another characterization models both the frequency and excesses simultaneously, which results in the use of the GEV distribution instead of the GP. In this case, the result is that the rate parameter is a function of the GEV parameters and the time interval. If excesses over a threshold follow a GP distribution, and occur at time points that follow a Poisson process, then maxima over disjoint intervals of fixed length follow a GEV distribution (see e.g., Coles 2001). A difficulty in estimating this characterization of the Poisson model from a software perspective is in estimating the rate parameter efficiently.

Many methods for estimating the parameters of EVDs are available. The most popular include: maximum likelihood (ML, e.g., Coles 2001), probability weighted moments (e.g., Hosking et al. 1985), L-moments (e.g., Hosking 1990), and Bayesian methods (e.g., Coles 2001; Cooley et al. 2007; Stephenson and Tawn 2004). Several estimators specific to the shape parameter have also been used (e.g., Beirlant et al. 2004; Hill 1975; Pickands 1975).

The ML problem is non-regular because the endpoints of the EVDs depend on the parameters, but it can be shown that the usual properties hold for the most typical situations whereby $\xi > -1/2$ (Smith 1985). Although real data rarely exhibit such short-tailed distributions, it has been noted that the likelihood can be made arbitrarily large when $\xi < -1$ and this prevents the use of the maximum likelihood estimator in such situations (Hosking and Wallis 1987). Because an analytic solution to maximizing the likelihood for the EVD's does not exist, numerical optimization is necessary. Nevertheless, the ML and Bayesian approaches lend themselves naturally to calculating uncertainty information and the extension of the EVD's to regression type equations. L-moments have desirable properties, such as providing good estimates in the face of small sample sizes, but do not extend easily to regression type equations. Further, uncertainty calculations are more involved.

For multivariate extremes, it is common to investigate componentwise maxima. That is, for a d -variate random vector, the block maxima of each component, under linear normalizations of each component, are analyzed. Thus, the multivariate aspect pertains to how the maxima of each individual component of the vector relate to each other. Although this assumption can be relaxed for concrete application, it is often more convenient to introduce the theory with arbitrarily fixed margins and a widely used choice is to assume unit Fréchet margins. This assumption will be assumed throughout the remainder of this section. The class of multivariate EVD's for componentwise maxima can be written as

$$G(z_1, \dots, z_d) = \exp\{-V(z_1, \dots, z_d)\}, \quad z_1 > 0, \dots, z_d > 0, \quad (2.2)$$

where

$$V(z_1, \dots, z_d) = \int_{S_d} \max_{j=1, \dots, d} \frac{w_j}{z_j} dH(w),$$

with H a positive measure on the unit simplex,

$$S_d = \left\{ (w_1, \dots, w_d) \in [0, \infty)^d \setminus \{\mathbf{0}\} : \sum_{j=1}^d w_j = 1 \right\},$$

and satisfying the constraints $\int_{S_d} w_j dH(w) = 1, j = 1, \dots, d$ to ensure unit Fréchet margins. Similar theory exists for threshold excesses (see, e.g., Coles and Tawn 1991; Ledford and Tawn 1996; Rootzén and Tajvidi 2006). For the class of multivariate EVDs, each bivariate margin yields either exact independence or asymptotic dependence (Coles et al. 1999). The expansion of models beyond this class to incorporate asymptotic independence has been investigated by Ledford and Tawn (1997), Coles and Pauli (2002) and Apputhurai and Stephenson (2011).

Max-stable processes are a key concept in analyzing extremes of stochastic processes. In accordance with the multivariate case, max-stable processes arise as the point-wise maxima of independent replicates of a continuous sample path stochastic process—again under linear normalizations. Interestingly, max-stable processes can be built through their spectral characterizations (e.g. de Haan 1984; Schlather 2002) one of which being

$$Z(x) = \max_{i \geq 1} \psi_i Y_i(x), \quad x \in \mathbb{R}^d, \tag{2.3}$$

where $\{\psi_i\}_{i \geq 1}$ are the points of a Poisson process on $(0, \infty)$ with intensity $d\Lambda(\psi) = \psi^{-2}d\psi$ and Y_i are independent replicates of a non negative with continuous sample path stochastic process Y such that $\mathbb{E}[Y(x)] = 1$ for all $x \in \mathbb{R}^d$. Different choices for the process Y lead to some useful max-stable models such as Smith (Smith 1990), Schlather (Schlather 2002) and Brown–Resnick (Kablichko et al. 2009) processes. Davison et al. (2011) give a detailed account of these models and compare them on the areal modeling of extreme rainfall events in Switzerland. Based on Eq. 2.3, it is not difficult to show that the multivariate distribution is

$$\Pr[Z(x_1) \leq z_1, \dots, Z(x_k) \leq z_k] = \exp \left\{ -\mathbb{E} \left[\max_{j=1, \dots, k} \frac{Y(x_j)}{z_j} \right] \right\}, \quad x_1, \dots, x_k \in \mathbb{R}^d,$$

and that, as expected, the finite dimensional distribution of a max-stable process has the same structure as the multivariate EVD.

Regional frequency analysis is a branch of extreme-value analysis in which multivariate data are available but the main focus is on estimation of extreme quantiles of the marginal distributions, with relatively little importance given to estimating the frequency of events that involve simultaneous extremes in multiple variates. This is appropriate for the design of buildings, dams, and other structures that are required to withstand extreme events of specified return periods. In this case the variates correspond to environmental quantities observed at different measuring sites. A common approach is to base frequency estimates on summary statistics such as moments or

L -moments computed separately at each site and pooled across the sites to give a single set of summary statistics applicable across a geographic area. This approach was proposed by Dalrymple (1960) and elaborated by, among others, Wallis (1980) and Hosking and Wallis (1997). Regional frequency analysis does not require the use of extreme-value distributions, but because the data studied are often environmental extremes, the use of extreme-value distributions in regional frequency analysis is widespread.

3 Current software

Focus here is on software packages available in the free, open-source statistical software language and environment called R (R Development Core Team 2012) because this has become the software language most used by academic statisticians, and because it is open-source and freely available without proprietary licensing requirements. The source code is therefore available for any user to: (i) help them understand what the software is actually doing, (ii) potentially find bugs, and in some cases, (iii) improve the code and make the modified version available for all to use. Further, R is the software that contains the most utilities for modeling extreme values. However, some detail is given for non R packages known to the authors at the time of writing.

A couple of the R packages we will discuss were originally written for the similar commercial software environment called S-Plus. For example, the R package `evir` was ported from the former S-Plus library `EVIS`, which has now been subsumed into the S-Plus module `S+FinMetrics`. All R packages we mention are currently free of cost. All packages are currently available from CRAN with the exception of `rgarch` and `inla`. To our knowledge, all packages are released under the GNU General Public Licence unless otherwise stated.

For general help with an R package, typing `help(package="name")` where `name` is the package name will give package information and a list of online help topics for that package. The package information includes the name and email address of the current maintainer, to whom all queries about the package should be directed. The online help topics can be accessed directly by typing `?topic` where `topic` is the topic name. The packages must be installed before accessing the online help in this manner. A package can be installed using `install.packages("name")`, which will automatically download it from CRAN (the comprehensive R archive network) and install it. Any other documentation available will be located in the `doc` directory underneath wherever the package is installed. To cite a package in publications please use the result returned by typing `citation("name")`.

3.1 R software packages

Table 1 provides a rough summary of the techniques implemented by each of the R software packages. More thorough details of each are provided in the text below, given in alphabetical order.

Table 1 General summary of R software packages for extreme value analysis

	Block maxima	Threshold models	Estimation methods	Non-stationary regression	Multivariate capability
copula	No	No	MLE, pseudo-MLE, MOM	No	Yes
evd	Yes	Yes	MLE	Some	Bivariate
evdbayes	Yes	Yes	Bayesian	Limited	No
evir	Yes	Yes	MLE	No	Limited
extRemes/ismev	Yes	Yes	MLE, LM	Yes	No
fExtremes	Yes	Yes	MLE, PWM	No	No
lmom	Yes	Yes	LM	No	No
lmomRFA	Yes	Yes	LM	No	Limited
lmomco	Yes	Yes	LM	No	Yes
POT	No	Yes	PMLE, Other	No	Bivariate
SpatialExtremes	Yes	Yes	MLE, MCLE, Bayesian	Yes	Yes
texmex	No	Yes	MLE, PMLE, Bayesian	Yes	Yes
VGAM	Yes	Yes	MLE, BFA	Yes	No

PMLE = penalized maximum likelihood estimation, *MLE* = maximum likelihood estimation, *MCLE* = maximum composite likelihood estimation, *MOM* = method of moments, *PWM* = probability weighted moments estimation, *LM* = L-moments estimation, *BFA* = backfitting algorithm

Copula (Kojadinovic and Yan 2010) A copula is simply a multivariate distribution whose univariate margins are uniformly distributed on the interval $[0, 1]$. If F is a d -dimensional distribution function with univariate marginals F_i for $i = 1, \dots, d$, then there exists a corresponding copula C such that

$$F(z_1, \dots, z_d) = C\{F_1(z_1), \dots, F_d(z_d)\},$$

and therefore C can be viewed as containing information on the dependence structure of F . If the margins F_i are known, then the distribution is defined only by C . In most applications, copulas are used under the assumption that the marginal distributions are continuous, which ensures that C is unique. When the margins are not continuous, the copula is unique only on the range of the random variables transformed by the marginal cumulative distribution functions.

The copula package contains utilities for exploring and modeling a wide range of commonly used copulas. This includes two elliptical copulas (F being a multivariate normal or multivariate t distribution), five extreme value copulas (F being a parametric subclass of the class of bivariate extreme value distributions), three Archimedean copulas and two other copulas (Plackett and Farlie–Gumbel–Morgenstern) that do not belong in the previous families. The package allows easy calculation of dependence measures based on copulas such as Kendall's τ , Spearman's ρ and the upper tail dependence index defined as the limit of $\Pr(F(z_1) > u | F(z_2) > u)$ as $u \uparrow 1$.

Non-parametric estimates of extreme value copulas are implemented, and various hypothesis tests can be performed, such as tests of extreme-value dependence and tests for goodness-of-fit. Algebraic expressions for the upper tail dependence index and other measures of extremal dependence are available in Heffernan (2000) for a large number of copula families. The upper tail dependence index is denoted by c in that reference, and should not be confused with the coefficient of tail dependence η .

evd (Stephenson 2002) The *evd* package was originally designed to provide additional functionality for distributions often encountered in the analysis of extreme values, but has subsequently been extended to include functions for statistical modeling using maximum likelihood estimation. A user's guide is available in the `doc` directory. Functions are provided for simulation, and for calculation of distribution, density and quantile functions, for a wide range of univariate and multivariate distributions. Cluster identification and extremal index estimation are also included.

Modeling functions allow estimation of parameters for standard univariate extreme value methods, and estimation for the class of bivariate extreme value distributions. Parametric and non-parametric estimation of bivariate extreme value distributions can be performed: parametric estimation can be performed for eight commonly used parametric sub-classes, and non-parametric estimation can be performed using the estimation methods of Pickands (1981) and Capéreaux et al. (1997). Graphical model diagnostics are provided for all fitted models. Profile likelihoods and profile likelihood confidence intervals can be calculated for any parameter of interest, including return levels. Bivariate threshold modeling is implemented using censored likelihood methodology (Coles 2001).

evdbayes (Stephenson and Ribatet 2010) The *evdbayes* provides functions for the Bayesian analysis of univariate extreme value models using Markov chain Monte Carlo (MCMC) methods. The utility of the package is restricted to univariate methods employing likelihoods for component-wise maxima, order statistics and threshold modeling. The package produces Markov chain output; MCMC diagnostics for features such as the burn-in period or chain length can then be produced using the R package `coda`.

The focus of the package is on informative Bayesian analysis, and therefore there are several functions employing different methods for the construction of informative priors over the relevant parameter space. The central feature of interest in a univariate Bayesian analysis of extremes is typically the posterior predictive distribution of return levels, and this can easily be plotted and calculated within the package. Three worked examples are available from the users guide, located in the `doc` directory.

evir (McNeil and Stephenson 2008) The package *evir* is primarily designed for applying extreme value methodology to financial data. It performs standard stationary univariate extreme value modeling, including maximum likelihood fitting of generalized extreme and generalized Pareto distributions. Functions exist for calculating expected shortfalls as well as quantiles, for extracting records and declustering, and for estimating the extremal index. It can plot the Hill estimate of the tail index

for heavy-tailed data, and has a function for extracting records. The single multivariate implementation is a bivariate threshold method, but the range of dependence structures is severely limited and more flexible options exist in the `evd` and `POT` packages.

extRemes and *ismev* (Gilleland and Katz 2011; Stephenson 2011) The package *ismev* was ported into R by Alec G. Stephenson from the original S-plus functions written by Janet E. Heffernan to accompany the text Coles (2001). An S-plus version of the *ismev* package (with the same name) is also available from the Comprehensive S-Plus Archive Network (CSAN, <http://csan.insightful.com>; Southworth 2007). It provides for univariate estimation of EVD's for the block maxima and threshold model approaches (including both GPD and Poisson Point Process models) obtained by MLE. Of note, it allows for incorporation of non-stationarity through the parameters of the EVD's of the form

$$f_{\phi} \{ \phi(x_1, \dots, x_m) \} = \phi_0 + \phi_1 x_1 + \dots + \phi_m x_m,$$

where ϕ can be the location, scale and/or shape parameter, f is a link function (e.g., $f(\sigma) = \log(\sigma)$ to ensure that $\sigma(x) > 0$ everywhere), and x_1, \dots, x_m are vectors of covariates of interest (e.g., time). That is, it allows for linear regression type accounting for non-stationary, possibly including a link function. The package also includes functions for diagnosing the quality of the fitted distributions (e.g., probability and qq-plots, histograms) as well as functions for assisting in selecting an appropriate threshold for the threshold models.

The package *extRemes* is essentially a GUI interface to the package *ismev*, but it does include some additional functionality. In particular, it allows for L-moments estimation for the stationary case for the GEV and GP distributions. Further, it has a more user-friendly command-line function for calculation of return levels (`return.levels`) that allows the user to store the output in an R object (*ismev* provides a return level plot through the diagnostic plot function, `gev.diag`, for pre-determined return periods, and does not return the values for further use). It also has some capability added by Christopher A. T. Ferro for calculating the extremal index and number of clusters (cf. Ferro and Segers 2003).

fExtremes and *Rmetrics* (Wuertz 2009) There are 27 R packages that are made available under the collective term of *Rmetrics*, an open source solution for teaching financial market analysis and valuation of financial instruments, and *fExtremes* is one of them. The package was developed using code from other existing R packages, particularly *evd* and *evir* and consequently there is a large overlap in utility.

The available package functions are for univariate data processing and modeling. Data processing includes clustering, exceedance and block maxima identification, exploratory data analysis, and extremal index estimation. Univariate simulation and distribution functions are implemented, as is estimation of stationary models for the GEV and GPD using maximum likelihood and probability weighted moments. The quantile based financial measure known as value-at-risk can be calculated explicitly,

and plots of tail index estimates can be produced using the estimators of Hill (1975), Pickands (1975) and Dekkers et al. (1989).

Installing `fExtremes` will automatically install the package `fBasics`, containing functions for stable distributions, which arise as a limiting distribution for sums of independent and identically distributed random variables with possibly infinite variance. Stable distributions are often used to model heavy-tailed data. The `fBasics` package provides simulation, distribution, density and quantile functions for the stable distribution. It performs parameter estimation using maximum likelihood and a quantile method due to McCulloch (1986). Another package of interest in `Rmetrics` is `fCopulae`, which appears to have almost identical functionality as the `copula` package on which it is presumably based.

lmom and lmomRFA (Hosking 2009a, b) The `lmom` package, written and maintained by J.R.M. Hosking, has functions to fit each of thirteen probability distributions to data using the low-order L -moments (Hosking 1990). These distributions include some that are widely used in extreme-value analysis: Gumbel, generalized extreme-value, and generalized Pareto. The package also contains functions to evaluate the cdf's, quantile functions and L -moments of these distributions. The package can also draw an L -moment ratio diagram, a tool useful in deciding which distributions may give an adequate fit to a specified data set.

`lmomRFA` extends package `lmom` and implements all the major components of the methodology described in Hosking and Wallis (1997) for regional frequency analysis using L -moments. The package provides functions to compute a discordancy measure, used to assess whether individual sites in a set of sites contain unusual or outlying data values; a heterogeneity measure, used to assess whether a set of sites is "homogeneous", i.e. whether a single frequency distribution can be applied to all the sites; and a goodness-of-fit measure, used to assess the adequacy of a specified distribution to fit the data from a homogeneous set of sites. The package can also fit a frequency distribution to a set of sites and construct estimates of the accuracy of the estimated quantiles for each site. Both `lmom` and `lmomRFA` are released under the Common Public License (CPL), which is less restrictive than GPL.

lmomCO (Asquith 2009) This package is similar to the aforementioned `lmom` package, but also implements several more recent advances in L -moment estimation, including L -moments for censored data, trimmed L -moments, and L -comoments for multivariate analysis.

POT (Ribatet 2009) The `POT` package, written and maintained by M. Ribatet, focuses only on the modeling of exceedances over a threshold. Multivariate exceedances are available but it is currently limited to the bivariate case only. This package differs from the other ones in that it provides many estimators for the GPD parameters ranging from the classical ones (MLE, L -moments / probability weighted moments) to robust ones (method of median, minimum density power divergence). It also provides additional techniques for threshold selection based on L -moments diagrams or from the properties of a homogeneous Poisson process. A users guide, provided as an R vignette, is accessible by invoking `vignette("POT")`.

SpatialExtremes (Ribatet 2011) The *SpatialExtremes* package, written and maintained by M. Ribatet, is devoted to the modeling of spatial extremes and provides several approaches for doing so. The first approach is based on max-stable processes and uses the (weighted) pairwise likelihood estimator to fit such processes to spatial data sets. Various parametric max-stable models can be used to describe the spatial dependence while trend surfaces, including temporal trends, can be used to model the spatial behavior of the marginal parameters. Simulation of several max-stable process models is also implemented as well as conditional simulation of (discretized) max-stable processes.

The second approach to modeling spatial extremes uses Bayesian hierarchical models with a conditional independence assumption. More precisely, it is assumed that the GEV parameters vary smoothly according to a stochastic process and that conditionally on this stochastic process the maxima $Y(x_j)$ satisfy

$$Y(x_j) \mid \{\mu(x_j), \sigma(x_j), \xi(x_j)\} \sim \text{GEV}\{\mu(x_j), \sigma(x_j), \xi(x_j)\}, \quad j = 1, \dots, k,$$

independently for each location (x_1, \dots, x_k) .

Copula approaches are also implemented (Davison et al. 2011). A user's guide, provided as an R vignette, is accessible by invoking `vignette("SpatialExtremesGuide")`.

texmex (Southworth and Heffernan 2010) The *texmex* package (Southworth and Heffernan 2010) implements the conditional multivariate extreme value modeling approach of Heffernan and Tawn (2004). This approach is useful for multivariate processes where interest is in events occurring such that only a subset of the margins are extreme. It also contains GPD modeling functionality using the R formula interface for each parameter. It has parametric bootstrap and MCMC simulation of posteriors, as well as both MLE and PMLE for parameter estimation.

Of particular note, *texmex* includes a large suite of test functions based on the RUnit package to perform several tests in order to ensure its functions behave as expected. These include tests to make sure that functions can reproduce published results (e.g., parameter estimates and model summaries that appear in Coles (2001) and Heffernan and Tawn (2004)), and many others (see the help file for `validate.texmex` for more information).

VGAM (Yee 2010) Vector generalized linear and additive models (Yee and Wild 1996), as implemented by *VGAM*, offer a unifying framework for modeling parameters as linear or smooth functions of covariates. Extreme value models were recently added to *VGAM*. Yee and Stephenson (2007) present the details of such models and the method of estimation, and illustrate implementations of the software for extreme value analysis. Estimation for vector generalized linear models is performed using iteratively reweighted least squares, while estimation for vector generalized additive models is performed using a modified backfitting algorithm, employing a penalized likelihood for the smoothing splines. *VGAM* is the only package known to the authors that offers additive modeling for a range of extreme value analyses.

The models available can be specified using the `family` argument of the `vgam` function. In terms of extreme value analysis, current implementation is focused on univariate models, however experienced users can construct their own family functions in order to extend the modeling capabilities. The package is very powerful and flexible, and models can be constructed that are unavailable in other software implementations; the disadvantage of such modeling flexibility is that the onus is on the user to specify models carefully so that they remain estimable.

Other R software In addition to the R packages listed above, other packages exist which may be useful for extremes or for which functions related to extremes comprise a small part. The package `actuar` provides functions related to different forms of Burr and Pareto distributions, as does the package `Runuran` for simulation. The packages `nacopula` and `fgac` provide some utilities for copula modeling. The package `gamlss` is for generalized additive modeling of distributions with location, scale and shape parameters, and provides an alternative to VGAM for univariate block maxima models. Conditional simulation of (discretized) max-stable processes can be performed using the `maxLinear` package; further utilities for simulating and analyzing max-stable processes are provided by `RandomFields`. Quantile regression facilities are provided by `quantreg`.

For volatility modeling with GARCH processes see `tseries` and `fGarch`. Also see `ccgarch` for condition correlation models and `gogarch` for generalized orthogonal models. The `rgarch` package aims to provide a flexible modeling and testing environment for garch models, but at the time of writing it was in beta testing and is not yet available from CRAN. The package `QRMLib` contains functions for quantitative risk management: it includes utilities for stationary univariate extreme value models and for copula modeling. `CreditMetrics` provides a model for credit risk. The package `inla` estimates a class of models using a novel inferential approach developed by Rue et al. (2009). This class includes models based on the generalized extreme value distribution. The package is not currently on CRAN but can be downloaded from <http://www.r-inla.org/>. Most of the code in `inla` is made available under GPL, but some external libraries have different licences.

The `extremevalues` package (van der Loo 2010), employs distribution-based outlier detection, along with plotting routines and a graphical user interface. In detecting outliers based on a distribution, the user can choose from: lognormal, exponential, Pareto, Weibull and normal. The package does not fit EVDs to data *per se*. Instead, log-linear regression is used to fit data values between two quantiles to the cdf of a model distribution, some of which are EVDs.

Although not available on CRAN at the time of writing, the package `TestEVC1d.r` includes functionality for testing the extreme value conditions (Hüsler and Li 2006a). In particular, the methods introduced in Dietrich et al. (2002) and Drees et al. (2006), and compared in Hüsler and Li (2006b), can be employed.

3.2 Other EVA software

Of course, other software besides the numerous packages in R are also available. Here, we briefly discuss some of the main ones known to the authors at the time

of writing. One that stands out is a stand-alone GUI package accompanying the text Reiss and Thomas (2007), *Xtremes*. A professional version of the package, which allows for large datasets to be analyzed, is available for purchase at <http://www.risktec.de/software.htm>. *Xtremes* represents a fairly complete analysis package aimed primarily at finance and hydrology applications, but the methods are generally applicable. It also includes the package *STABLE* written by John P. Nolan (see e.g. Nolan 2007).

A free MatLab package for EVA is the *EVIM* package (Gençay et al. 2001). It is currently available as a zip file from <http://www.sfu.ca/~rgenay/>, and contains univariate routines that are similar to the R package *evir*, such as those for block maxima and peaks over threshold data, declustering methods, hill plots and extremal index estimation. Unfortunately it no longer appears to be in development. More freely available MatLab routines for standard univariate methods are provided by *WAF0* (Brodtkorb et al. 2000), which is also available in the Python language, in the context of wave analysis for oceanography. It also contains routines for sea modeling and fatigue analysis. Version 2.5 of *WAF0* was recently released, and contains many new routines including those for profile log-likelihoods and bootstrap confidence interval estimation, as well as the maximum product of spacings (MPS) method (Wong and Li 2006).

EXTREMES is a GUI software package written in C++ that provides a wealth of functionality for univariate extreme value analysis including block maxima and threshold models (Diebolt et al. 2003a). It also includes many standard statistical methods (e.g., sample simulation, parameter estimation, non-parametric density estimation, parametric quantile estimation, Anderson-Darling test, etc.). In order to check that a parametric model provides acceptable tail approximations, *EXTREMES* includes a test which compares the parametric estimate of an extreme quantile with its semiparametric estimate obtained by extreme value theory (Diebolt et al. 2007). The sampling variations of these estimates are approximated through parametric bootstrap. *EXTREMES* also proposes a regularization procedure which preserves the general form of a parametric distribution and allows a better fit in the distribution tail (Diebolt et al. 2003b). It is based on Bayesian tools and takes the opinions of experts into account. Predictive distributions are proposed as model distributions. They are obtained as a mixture of the model family density functions according to the posterior distribution. The software can be downloaded at <http://extremes.gforge.inria.fr/>.

HYFRAN and *HYFRAN-PLUS* are commercial standalone software products (El Adlouni et al. 2008) authored by a group of researchers at INRS University, Canada, originally designed for hydrological frequency analysis but now marketed more generally as software for fitting a large number of distributions with different tail weights to *iid* data using a wide range of estimation methods. *HYFRAN-PLUS* contains what they call a Decision Support System (DSS), which includes graphical diagnostics to enable the user to choose between different models. Both products can be purchased from <http://www.wrpllc.com/books/hyfran.html>. The main advantage of *HYFRAN* and *HYFRAN-PLUS* is the provision of an easy-to-use interface for those with little computational experience. Free demo versions are available, but these only allow estimation using the normal distribution. WINFAP-FEH 3

is another commercial standalone hydrological software product, distributed by Wallingford HydroSolutions. It implements flood frequency analysis techniques using both annual maxima and peaks over threshold methods.

`S+FinMetrics` is an S-Plus module for econometric analysis. An S-Plus module must be purchased in addition to S-Plus software. It contains tools for modeling and analyzing market data. This includes tools for time series analysis, econometric system estimation, and state space modeling. It also provides tools for volatility modeling such as GARCH models, for copula modeling and estimation, and for univariate extreme value methodology. The code for the extreme value methodology derives from the same code used for the R package `evir`, and so contains similar utilities. The set of `Rmetrics` R packages has been designed to provide a free alternative to `S+FinMetrics`, in particular `fExtremes`, `fCopulae` and `fGarch` cover similar ground.

We conclude with two stand-alone software packages written in Fortran and freely available from the U.S. Geological Survey's Water Resources Applications Software (<http://water.usgs.gov/software/>). Both packages are devoted to the modeling of flow characteristics but do not include any extreme value distributions. The first one is `GLSNet`, which performs prediction of low/high flows at non-instrumented locations using regression techniques and generalized least squares. The second package, `peakFQ`, estimates annual peak flows for several return periods by fitting a Pearson Type III distribution, i.e., a Gamma distribution, using the logarithmic sample moments.

4 Discussion

Since Stephenson and Gilleland (2005), new development of software for the analysis of extreme values has been driven primarily by the creation of new R packages from individual researchers. This will inevitably continue into the future, with an ever expanding set of R packages containing routines for fitting statistical extreme event models. This is helped by the fact that academic recognition for software production is increasing, with several outlets now available for publication of this important but time-consuming work.

Perhaps the most valuable software developed since 2005 is that which relates to spatial modeling. R packages such as `SpatialExtremes` allow the user to perform inference on spatial extremes relatively easily for even complex models, providing valuable functionality that was not available five years ago. The applicability of R software is now limited more by theoretical constraints than by the availability of algorithms. For example, the lack of availability of efficient simulation algorithms for certain classes of max-stable processes (e.g. Oesting et al. 2012) is an important theoretical constraint for `SpatialExtremes`.

Stephenson and Gilleland (2005) proposed the creation of a software initiative to develop a reliable and coherent set of tools in order to further increase the use of extreme value methods in various applied fields. Since that time, however, there has not been an identifiable collaborative framework; progression has continued largely through the efforts of individuals. The advantage of this approach is that development

has been rapid, while the disadvantage is that there are now more packages with similar routines, creating confusion for the end-user.

Some packages are approaching a global goal (e.g., `SpatialExtremes`), and other more specialized packages (e.g., `lmom`, `evd`, `evir`, `extRemes`, `ismev`, etc.) can be used together to perform most extreme value analyses. It would be useful to make some of these more consistent with each other so that together they could make a comprehensive set of tools that can readily be used or adapted by professionals needing to incorporate them in their work. However this would require both a large amount of collaboration and potentially large changes to individual packages.

The expanding set of R packages provides complexity for the end user with regard to understanding what packages are currently available and what each package offers, particularly given the rapid pace of development. This issue is not unique to extreme value theory, and an attempt to summarize available packages in R relating to particular statistical topics has been made through so-called task views, which allow the user to browse packages by topic and provide tools to automatically install all packages for a particular area of interest. No task view currently exists for extreme value theory.

Packages in R have reasonably good documentation for individual functions as the R developers carefully ensured through the process by which such packages are uploaded to CRAN. The checking tools for R developers are far more thorough than they once were. Some packages (e.g., `SpatialExtremes`, `extRemes`) have tutorial documentation describing how to use the packages to perform useful analyses. However because of the nature of development, there is a lack of documentation that cuts across specific software units.

The development of non R software for extremes has not been so rapid. In particular, to our knowledge there has been no new commercial software product incorporating extreme value analysis since Stephenson and Gilleland (2005). The advantage of such products is primarily ease-of-use for practitioners that may be discouraged by the learning curve that R presents. It is not surprising then that existing products are typically aimed at practitioners in a particular area of application: `S+FinMetrics` for finance, `HYFRAN-PLUS` for hydrology and the freely available `WAFO` for oceanography. The main commercial product available is still `S+FinMetrics`, which is marketed toward companies in the business sector. The lack of commercial competition leads to the interesting question of the extent to which practitioners are using extreme value methods, given that R is largely popular in academia.

Acknowledgments We would like to warmly thank Stéphane Girard, Institut National de Recherche en Informatique et en Automatique, and Jonathan R. Hosking, IBM Research Division, for contributing the sections on the `EXTREMES` software and the `lmom` R package respectively. We also thank the four anonymous reviewers for their helpful suggestions that have made this a better paper.

Partial funding for this work was provided by the Weather and Climate Impacts Assessment Science Program of the National Center for Atmospheric Research, which is funded by the National Science Foundation. M. Ribatet was partly funded by the `MIRACCLE-GICC` and `McSim` ANR projects.

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

References

- Apputhurai, P., Stephenson, A.G.: Accounting for uncertainty in extremal dependence modeling using Bayesian model averaging techniques. *J. Stat. Plan. Inference* **141**, 1800–1807 (2011)
- Asquith, W.H.: *lmomco: L-moments, Trimmed L-moments, L-comoments, and Many Distributions*. R package version 0.97.4 ed. (2009)
- Beirlant, J., Goegebeur, Y., Segers, J., Teugels, J.: *Statistics of Extremes*. Wiley, Chichester (2004)
- Brodtkorb, P., Johannesson, P., Lindgren, G., Rychlik, I., Rydén, J., Sjö, E., WAFO—a Matlab toolbox for the analysis of random waves and loads. In: *Proc. 10th Int. Offshore and Polar Eng. Conf.*, vol. 3. ISOPE, Seattle, USA (2000)
- Capéraà, P., Fougères, A.-L., Genest, C.: A non-parametric estimation procedure for bivariate extreme value copulas. *Biometrika* **84**, 567–577 (1997)
- Coles, S.G.: *An Introduction to Statistical Modeling of Extreme Values*. Springer, London (2001)
- Coles, S., Pauli, F.: Models and inference for uncertainty in extremal dependence. *Biometrika* **89**, 183–196 (2002)
- Coles, S.G., Tawn, J.A.: Modelling extreme multivariate events. *J. R. Stat. Soc. B* **53**, 377–392 (1991)
- Coles, S.G., Heffernan, J.E., Tawn, J.A.: Dependence measures for extreme value analyses. *Extremes* **2**, 339–365 (1999)
- Cooley, D., Nychka, D.W., Naveau, P.: Bayesian spatial modeling of extreme precipitation return levels. *J. Am. Stat. Assoc.* **102**, 824–840 (2007)
- Dalrymple, T.: *Flood frequency analyses*. Water Supply Paper 1543-A, U.S. Geological Survey, Reston, VA (1960)
- Davison, A.C., Padoan, S.A., Ribatet, M.: Statistical modelling of spatial extremes. *Stat. Sci.* **27**(2), 161–186 (2012)
- de Haan, L.: A spectral representation for max-stable processes. *Ann. Probab.* **12**, 1194–1204 (1984)
- de Haan, L., Ferreira, A.: *Extreme value theory: an introduction*. In: *Springer Series in Operations Research and Financial Engineering*, 418pp. Springer, New York (2006)
- Dekkers, A.L.M., Einmahl, J.H.J., de Haan, L.: A moment estimator for the index of an extreme-value distribution. *Ann. Stat.* **17**, 1833–1855 (1989)
- Diebolt, J., Ecarnot, J., Garrido, M., Girard, S., Lagrange, D.: Le logiciel Extremes, un outil pour l'étude des queues de distribution. *Revue Modulad* **30**, 53–60 (2003a)
- Diebolt, J., Garrido, M., Trottier, C.: Improving extremal fit: a Bayesian regularization procedure. *Reliab. Eng. Syst. Saf.* **82**(1), 21–31 (2003b)
- Diebolt, J., Garrido, M., Girard, S.: A goodness-of-fit test for the distribution tail. In: Ahsanullah, M., Kirmani, S. (eds.) *Extreme Value Distributions*, pp. 95–109. Nova Science, New York (2007)
- Dietrich, D., de Haan, L., Hüslér, J.: Testing extreme value conditions. *Extremes* **5**, 71–85 (2002)
- Drees, H., de Haan, L., Li, D.: Approximations to the tail empirical distribution function with application to testing extreme value conditions. *J. Stat. Plan. Inference* **136**, 3498–3538 (2006)
- El Adlouni, S., Bobée, B., Ouarda, T.B.M.J.: On the tails of extreme event distributions in hydrology. *J. Hydrol.* **355**, 16–33 (2008)
- Embrechts, P., Klüppelberg, C., Mikosch, T.: *Modelling Extremal Events for Insurance and Finance*, 648pp. Springer, Berlin (1997)
- Ferro, C.A.T., Segers, J.: Inference for clusters of extreme values. *J. R. Stat. Soc. B* **65**, 545–556 (2003)
- Gençay, R., Selçuk, F., Ulugülyağci, A.: EVIM: a software package for extreme value analysis in MATLAB. *Stud. Nonlinear Dyn. Econom.* **5**(3), 213–239 (2001)
- Gilleland, E., Katz, R.W.: New software to analyze how extremes change over time. *Eos* **92**(2), 13–14 (2011)
- Heffernan, J.E.: A directory of coefficients of tail dependence. *Extremes* **3**, 279–290 (2000)
- Heffernan, J.E., Tawn, J.A.: A conditional approach for multivariate extreme values (with discussion). *J. R. Stat. Soc., Ser. B* **66**, 497–546 (2004)
- Hill, B.M.: A simple general approach to inference about the tail of a distribution. *Ann. Stat.* **3**, 1163–1174 (1975)
- Hosking, J.R.M.: L-moments: analysis and estimation of distributions using linear combinations of order statistics. *J. R. Stat. Soc., Ser. B* **52**, 105–124 (1990)
- Hosking, J.R.M.: *L-moments*, R package version 1.5 ed. (2009a)
- Hosking, J.R.M.: *Regional frequency analysis using L-moments*, R package version 2.2 ed. (2009b)
- Hosking, J.R.M., Wallis, J.R.: Parameter and quantile estimation for the Generalized Pareto distribution. *Technometrics* **29**(3), 339–349 (1987)

- Hosking, J.R.M., Wallis, J.R.: *Regional Frequency Analysis: An Approach Based on L-Moments*. Cambridge University Press, Cambridge (1997)
- Hosking, J.R.M., Wallis, J.R., Wood, E.F.: Estimation of the generalized extreme-value distribution by the method of probability-weighted moments. *Technometrics* **27**, 251–261 (1985)
- Hüsler, J., Li, D.: How to use the package TestEVC1d.r, 3pp. Available at: <http://my.gi.fudan.edu.cn/teacherhome/lideyuan/research.html> (2006a)
- Hüsler, J., Li, D.: On testing extreme value conditions. *Extremes* **9**, 69–86 (2006b)
- Kabluchko, Z., Schlather, M., de Haan, L.: Stationary max-stable fields associated to negative definite functions. *Ann. Probab.* **37**(5), 2042–2065 (2009)
- Kojadinovic, I., Yan, J.: Modeling multivariate distributions with continuous margins using the copula R package. *Journal of Statistical Software* **34**, 1–20 (2010)
- Ledford, A.W., Tawn, J.A.: Statistics for near independence in multivariate extreme values. *Biometrika* **83**, 169–187 (1996)
- Ledford, W.A., Tawn, J.A.: Modelling dependence within joint tail regions. *J. R. Stat. Soc. B* **59**, 475–499 (1997)
- McCulloch, J.H.: Simple consistent estimators of stable distribution parameters. *Commun. Stat., Simul. Comput.* **15**, 1109–1136 (1986)
- McNeil, A., Stephenson, A.G.: *evir: extreme values in R* (2008)
- Nolan, J.P.: *Stable Distributions—Models for Heavy Tailed Data*, 352pp. Birkhauser, Boston (2007). ISBN-13: 9780817641597
- Oesting, J., Kabluchko, Z., Schlather, M.: Simulation of Brown–Resnick processes. *Extremes* **15**(1), 89–107 (2012). doi:10.1007/s10687-011-0128-8
- Pickands, J.: Statistical inference using extreme order statistics. *Ann. Stat.* **3**, 119–131 (1975)
- Pickands, J.: Multivariate extreme value distributions. In: *Proc. 43rd Sess. Int. Statist. Inst.*, vol. 49, pp. 859–878 (1981)
- R Development Core Team: *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria (2012). ISBN 3-900051-07-0
- Reiss, R.D., Thomas, M.: *Statistical Analysis of Extreme Values, From Insurance, Finance Hydrology and Other Fields*. Birkhauser, New York (2001)
- Reiss, R.D., Thomas, M.: *Statistical Analysis of Extreme Values with Applications to Insurance, Finance, Hydrology and Other Fields*, 3rd edn. Birkhauser, New York (2007)
- Ribatet, M.: *POT: Generalized Pareto Distribution and Peaks Over Threshold*, R package verions 1.1-0 ed. (2009)
- Ribatet, M.: *SpatialExtremes: Modelling Spatial Extremes*, R package version 1.8-5 (2011)
- Rootzén, H., Tajvidi, N.: Multivariate generalized Pareto distributions. *Bernoulli* **12**(5), 917–930 (2006)
- Rue, H., Martino, S., Chopin, N.: Approximate Bayesian inference for latent Gaussian models using integrated nested Laplace approximations (with discussion). *J. R. Stat. Soc. B* **71**, 319–392 (2009)
- Schlather, M.: Models for stationary max-stable random fields. *Extremes* **5**(1), 33–44 (2002)
- Smith, R.L.: Maximum likelihood estimation in a class of non-regular cases. *Biometrika* **72**, 67–90 (1985)
- Smith, R.L.: Max-stable processes and spatial extreme. <http://www.stat.unc.edu/postscript/rs/spatex.pdf> (1990)
- Southworth, H.: *ismev: An Introduction to Statistical Modeling of Extreme Values, Original S functions written by Janet E. Heffernan, S-PLUS package by Harry Southworth*. S-PLUS package version 1.2 ed. (2007)
- Southworth, H., Heffernan, J.E.: *texmex: Threshold exceedences and multivariate extremes*, R package version 1.0 (2010)
- Stephenson, A.G.: *evd: extreme value distributions*. *R News* **2**(2), 31–32 (2002)
- Stephenson, A.G.: *ismev: An Introduction to Statistical Modeling of Extreme Values, Original S functions written by Janet E. Heffernan with R port and documentation provided by A. G. Stephenson*. R package version 1.35 ed. (2011)
- Stephenson, A.G., Gilleland, E.: Software for the analysis of extreme events: the current state and future directions. *Extremes* **8**, 87–109 (2005)
- Stephenson, A.G., Ribatet, M.: *evdbayes: Bayesian analysis in extreme value theory*, R package version 1.0-8 ed. (2010)
- Stephenson, A.G., Tawn, J.A.: Bayesian inference for extremes: accounting for the three extremal types. *Extremes* **7**, 291–307 (2004)

- van der Loo, M.P.J.: Distribution Based Outlier Detection for Univariate Data. Statistics Netherlands, The Hague (2010)
- Wallis, J.R.: Risk and uncertainties in the evaluation of flood events for the design of hydraulic structures. In: Guggino, E., Rossi, G., Todini, E. (eds.) *Piene e Siccità*, pp. 3–36. Fondazione Politecnica del Mediterraneo, Catania (1980)
- Wong, T.S.T., Li, W.K.: A note on the estimation of extreme value distributions using maximum product of spacings. *IMS Lecture Notes* **52**, 272–283 (2006)
- Wuertz, D.: *fExtremes: Rmetrics—Extreme Financial Market Data*, R package version 2100.77 ed. (2009)
- Yee, T.W.: The VGAM package for categorical data analysis. *Journal of Statistical Software* **32**, 1–34 (2010)
- Yee, T.W., Stephenson, A.G.: Vector generalized linear and additive extreme value models. *Extremes* **10**, 1–19 (2007)
- Yee, T.W., Wild, C.J.: Vector generalized additive models. *J. R. Stat. Soc. B* **58**, 481–493 (1996)