



***A SELF-CALIBRATING METHOD FOR HEAVY TAILED
DATA MODELING. APPLICATION IN NEUROSCIENCE AND
FINANCE***

RESEARCH CENTER
NEHLA DEBBABI, MARIE KRATZ, MAMADOU MBOUP
ESSEC WORKING PAPER 1619

DECEMBER 2016

A self-calibrating method for heavy tailed data modeling. Application in neuroscience and finance

Nehla Debbabi ^{*†‡}, Marie Kratz [§], Mamadou Mboup [†]

December 12, 2016

Abstract

One of the main issues in the statistical literature of extremes concerns the tail index estimation, closely linked to the determination of a threshold above which a Generalized Pareto Distribution (GPD) can be fitted. Approaches to this estimation may be classified into two classes, one using standard Peak Over Threshold (POT) methods, in which the threshold to estimate the tail is chosen graphically according to the problem, the other suggesting self-calibrating methods, where the threshold is algorithmically determined. Our approach belongs to this second class proposing a hybrid distribution for heavy tailed data modeling, which links a normal (or lognormal) distribution to a GPD via an exponential distribution that bridges the gap between mean and asymptotic behaviors. A new unsupervised algorithm is then developed for estimating the parameters of this model. The effectiveness of our self-calibrating method is studied in terms of goodness-of-fit on simulated data. Then, it is applied to real data from neuroscience and finance, respectively. A comparison with other more standard extreme approaches follows.

Keywords: Algorithm; Extreme Value Theory; Gaussian distribution; Generalized Pareto Distribution; Heavy tailed data; Hybrid model; Least squares optimization; Levenberg Marquardt algorithm; Neural data; S&P 500 index

2010 AMS classification: 60G70; 62E20; 62F35; 62P05; 62P10; 65D15; 68W40

JEL classification: C02

1 Introduction

Modeling non-homogeneous and multi-component data is a problem that challenges scientific researchers in several fields, as e.g. in climatology, finance & insurance, meteorology, neuroscience, ... (see e.g. [10, 26, 31, 32, 46, 49, 51]). In general, it is not possible to find a simple and closed form probabilistic model to describe such data. That is why one often resorts to non-parametric approaches, such as e.g. kernel density estimation ones (see e.g. [23, 40, 54]) or non-parametric Bayesian methods (see e.g. [2, 39, 53]), just to name a few. However, when the multiple components are separable, parametric modeling becomes again tractable. Several hybrid models have been proposed in such context, combining two or more densities (see e.g. [5, 25, 27, 31, 32]).

^{*}ESPRIT School of Engineering, Tunis, Tunisia, E-mail: nehla.debbabi@supcom.tn

[†]Université de Reims Champagne Ardenne, CReSTIC Lab., France, E-mail: mamadou.mboup@univ-reims.fr

[‡]University of Carthage, SUP'COM, COSIM Research Lab., Tunisia

[§]ESSEC Business School, CReAR risk research center; E-mail: kratz@essec.edu

In this study, we tackle the general problem in a specific case, when data exhibit heavy tails. Extreme behaviors that are described by heavy tail modeling, can be observed for a large number of phenomena, natural (from the big Dutch floods of 1952 to the recent earthquake in Italy), financial (e.g. the sub-prime crisis in North America or the Sovereign debt crisis in Europe), medical (e.g. the avian influenza), technological (e.g. Fukushima) or others. One mathematical field, Extreme Value Theory (EVT), which started with Tippet and Fisher (1928), is totally devoted to the analysis and modeling of the extremes (see [1, 8, 15, 20, 29, 47, 48] for general references). Studies on extremes were developed in many fields, as, for instance, in financial markets and actuarial mathematics (see e.g. [16, 18, 35]), in epidemiology (see e.g. [19] for the first introduction of EVT in epidemiologic methods), in signal processing (see e.g. [4, 17] when considering the general problem of false-alarms probability determination, or [14] for the spike detection in neural signals in biomedicine, or [36] for the detection of a binary signal in additive noise in telecommunication, or [56] for the damage detection in machine diagnostics).

Introducing EVT helps managing the many catastrophes that our society is facing with, unfortunately, an observed increasing trend of occurrence of extreme events since the beginning of the 20th century (see [6]), but also helps improving the standard data processing by taking into account the tail information.

Whereas EVT focuses on how to study and model extremes using the information in the tail of the distribution only (which is the strength of this theory, even if sometimes also its weakness in practice as tail data are scarce by definition), it is also very useful to combine it with standard statistics developed for the main information given in the data. To extract the important information given by extremes and to highlight as well the information contained in the entire underlying distribution, it is natural to take into account the dissymmetry of the data weights above a high threshold (tail) and below it (around the mean). Different methods have been proposed so far to do it (see e.g. [5, 25, 27, 31, 32, 37]).

The main goal of this paper is to develop a self-calibrating method to model heavy tailed data, the choice of this class of unsupervised procedures being clearly to ease practical implementations (in particular when complexity burden and/or delay processing are critical) and to enlarge its applicability. Indeed, the difficulty faced when applying standard methods of EVT as the Peaks Over Threshold (POT) approach (first introduced by Davison and Smith in the 90's; see [7]), the Hill method (see [21]), or the QQ-estimator one (see [28]), is that they are graphical ad hoc approaches.

This self-calibrating method may be seen as two-folds: when (i) looking for a full modeling for non-homogeneous, multi-component and heavy tailed data, (ii) focusing on the tail and evaluating in an unsupervised way the high threshold over which the tail will be modeled; it might then constitute an alternative EVT method to standard ones as e.g. the POT approach.

In this paper, we introduce a hybrid model with several components including a Generalized Pareto Distribution (GPD), to take into account the heavy tail present in the data. Without loss of generality, we assume continuous and asymmetric right heavy tailed data, a similar treatment being possible on the left tail (see [10]). How many components of the hybrid model to consider and how to choose them? Since we are interested in fitting the whole distribution underlying asymmetric heavy tailed data, the idea is to consider both the mean and tail behaviors, and to use limit theorems for each one (as suggested and developed analytically in [27]), in order to make the model as general

as possible. Therefore, we introduce a Gaussian distribution for the mean behavior, justified by the Central Limit Theorem (CLT), and a GPD for the tail, since the Pickands theorem (see [41]) tells us that the tail of the distribution may be evaluated through a GPD above a high threshold. To bridge the gap between mean and asymptotic behaviors, we use an exponential distribution. A different weight has been assigned to each component in order to have a better handling of the extremes. The resulting three-components hybrid model is called G-E-GPD model. Note that the GPD is the fixed component of this heavy tailed model, but the two other components could be chosen differently, depending on the data, and even reduced to one component (as developed earlier in [11,12]). Indeed, specific treatment could be done to fit the exact distribution of the mean behavior for which we have much data, if we would like to avoid the use of the limiting normal distribution. For instance, when having skewed distribution near the mean, which is typical for insurance claims data, the normal distribution should be replaced by a lognormal without loss of generality (see e.g. [25]). It would not change the idea of the self-calibrating method. Concerning the number of components, we point out that the model needs at least two-components, including the GPD, for the method to be workable. Indeed, the threshold over which the GPD is fitted (that we call the tail threshold), is determined in the algorithm as the junction point between the GPD and another distribution. Contrary to standard EVT approaches, it means that we need some information before the tail threshold to benefit from this self-calibrating method (further investigation will be made to adapt the method when partial information is available before the tail threshold).

Moreover, the intermediate distribution (here an exponential) is used as a leverage to give full meaning of tail threshold to the last junction point between the GPD and its neighbour (the intermediate distribution). The distance between two successive junction points will automatically tend to 0 when introducing unuseful components.

An iterative unsupervised algorithm is developed for estimating the parameters of the three-components hybrid model. It starts by enforcing the continuity and the differentiability of the three components at the two junction points, then proceeds in an iterative way to determine successive thresholds and parameters of the involved distributions. It provides a judicious weighting of the three distributions as well as a good location for the junction points or thresholds, especially for the tail threshold that points out the presence of extremes. This algorithm is based, for each iteration, on the resolution of numerical optimisation problems in least squares sense, using the Levenberg Marquardt (LM) method (e.g. [30,33]). We study its convergence analytically and numerically.

The performance of this self-calibrating method is studied in terms of goodness-of-fit on simulated data from G-E-GPD Monte-carlo simulations. Given the very good performance, we apply the method on real data, considering neural data and the *SP500* log-returns. A comparison with other existing graphical approaches is also given.

The paper is organized as follows. In section 2 we introduce our hybrid model. The method and its unsupervised iterative algorithm are developed in Section 3. Simulation results are presented in Section 4, and applications of the method on real data in Section 5. Results are discussed, in both sections, accompanied with a comparison of those obtained via standard methods. Conclusions follow in the last section.

2 Hybrid three-components model

We consider a piecewise model where each component represents a different behavior of the data, which might be heterogeneous or not. We assume that the data admit a continuous (non-degenerate) distribution, and accordingly, we introduce a general hybrid probability density function (pdf), with some smoothness constraints. Without any loss of generality, we consider a three-components model. More precisely, the hybrid model we propose, links three different distributions to each other at two junction points, denoted by u_1 and u_2 : a Gaussian distribution to model the mean behavior of the data, a GPD to represent the tail and an exponential distribution to bridge the gap between these two behaviors. This model, denoted by G-E-GPD (Gaussian-Exponential-Generalized Pareto Distribution), is characterized by its pdf h expressed as:

$$h(x; \theta) = \begin{cases} \gamma_1 f(x; \mu, \sigma) & \text{if } x \leq u_1, \\ \gamma_2 e(x; \lambda) & \text{if } u_1 \leq x \leq u_2, \\ \gamma_3 g(x - u_2; \xi, \beta) & \text{if } x \geq u_2, \end{cases}$$

The different parameters are gathered in the vector denoted by θ and are described hereafter. To begin, γ_i , $i = 1, 2, 3$ stand for the weights associated to each component. The parameters $\mu \in \mathbb{R}$, and $\sigma \in \mathbb{R}_+^* = \mathbb{R}_+ \setminus \{0\}$ represent, respectively, the mean and the standard deviation of the Gaussian pdf f given by: $f(x; \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$, $\forall x \in \mathbb{R}$. The parameters $\xi \in \mathbb{R}$ and $\beta \in \mathbb{R}_+^*$ denote, respectively, the tail index and the shape parameter of the GPD pdf g , defined by:

$$g(x; \xi, \beta) = \begin{cases} \frac{1}{\beta} (1 + \xi x/\beta)^{-1-\frac{1}{\xi}} & \text{if } \xi \neq 0 \\ \frac{1}{\beta} e^{-\frac{x}{\beta}} & \text{if } \xi = 0 \end{cases}, \quad \forall x \in \mathcal{D}(\xi, \beta),$$

where

$$\mathcal{D}(\xi, \beta) = \begin{cases} [0, \infty) & \text{if } \xi \geq 0 \\ [0, -\frac{\beta}{\xi}] & \text{if } \xi < 0 \end{cases}.$$

Finally, $\lambda \in \mathbb{R}_+^*$ indicates the intensity parameter of the exponential pdf e defined by $e(x; \lambda) = \lambda e^{-\lambda x}$, $\forall x > 0$.

In the sequel, we consider that the transitions from one behavior to another are smooth. As a consequence, we constraint the resulting hybrid pdf h to be \mathcal{C}^1 -regular. Note that by combining this constraint and the assumption of heavy tailed data, the number of free parameters and hence the size of θ will be reduced. Let us present these assumptions.

Assumptions of the model

The first two assumptions are part of the construction of the G-E-GPD model.

- (i) First we assume, by construction, that the data distribution admits a pdf h . This means that h is non-negative and satisfies $\int_{\mathbb{R}} h(x; \theta) dx = 1$, *i.e.*

$$\gamma_1 F(u_1; \mu, \sigma) + \gamma_2 (e^{-\lambda u_1} - e^{-\lambda u_2}) + \gamma_3 = 1, \quad (2.1)$$

where F denotes the cumulative distribution function (cdf) of the Gaussian distribution.

- (ii) We focus on heavy tailed data. This implies that h belongs to the Fréchet maximum domain of attraction ($\xi > 0$) and therefore $\beta = \xi u_2$ (see *e.g.* [15], p. 159).

The main constraint is to require a smooth pdf and to further reduce the number of free-parameters, that is why we have imposed h to be of class \mathcal{C}^1 ;

- (iii) h is continuous and differentiable at the two junctions points u_1 and u_2 .

Assumptions (i)-(iii) give rise to six equations relating all model parameters:

$$\begin{cases} \beta = \xi u_2; & \gamma_1 = \gamma_2 \frac{e(u_1; \lambda)}{f(u_1; \mu, \sigma)}; \\ \lambda = \frac{1+\xi}{\beta}; & \gamma_2 = \left[(\lambda\beta - 1) e^{-\lambda u_2} + \left(1 + \lambda \frac{F(u_1; \mu, \sigma)}{f(u_1; \mu, \sigma)} \right) e^{-\lambda u_1} \right]^{-1}; \\ u_1 = \mu + \lambda \sigma^2; & \gamma_3 = \beta \gamma_2 e(u_2; \lambda). \end{cases} \quad (2.2)$$

Consequently, the parameter vector θ retains only the free parameters and we set $\theta = [\mu, \sigma, u_2, \xi]$. It is then straightforward to deduce from h the expression of the cdf and quantile function associated with the G-E-GPD model. The G-E-GPD cdf, denoted H , is given by:

$$H(x; \theta) = \begin{cases} \gamma_1 F(x; \mu, \sigma) & \text{if } x \leq u_1, \\ \gamma_1 F(u_1; \mu, \sigma) + \gamma_2 (e^{-\lambda u_1} - e^{-\lambda x}) & \text{if } u_1 \leq x \leq u_2, \\ 1 - \gamma_3 (1 + \frac{\xi}{\beta} (x - u_2))^{-\frac{1}{\xi}} & \text{if } x \geq u_2, \end{cases} \quad (2.3)$$

and the corresponding quantile function by:

$$H^{-1}(p; \theta) = \begin{cases} F^{-1}\left(\frac{p}{\gamma_1}; \mu, \sigma\right) & \text{if } p \leq p_1 := \gamma_1 F(u_1; \mu, \sigma), \\ \lambda^{-1} \log\left(\frac{\gamma_2}{p_1 - p + \gamma_2 e^{-\lambda u_1}}\right) & \text{if } p_1 \leq p \leq p_2 := 1 - \gamma_3, \\ \frac{\beta}{\xi} \left(\left(1 - \frac{p - p_2}{\gamma_3} \right)^{-\xi} - 1 \right) + u_2 & \text{if } p \geq p_2, \end{cases}$$

where the notation Λ^{-1} denotes the inverse function of the function Λ .

A classical problem that arises when dealing with parametric models concerns how to estimate the model parameters. To answer this problem, we develop an iterative algorithm for estimating the parameters vector θ . This algorithm is an extension of the one built in [9, 11, 13]. For each iteration, it is based on the numerical resolution of optimization problems in least squares sense, using the Levenberg Marquardt (LM) method (see [30, 33]). We describe it and study its convergence in the next session.

3 Iterative algorithm for hybrid model parameters estimation

Here we describe the iterative algorithm suggested to estimate the G-E-GPD model parameters, which self-calibrates the model, in particular the threshold above which a Fréchet distribution fits the extremes. This algorithm follows the same logic as the one developed for two-components in [9, 11, 13]. We will recall it in the appendix when studying its convergence. For each iteration, it

breaks down the problem of the parameters vector θ estimation into two nested subproblems; the parameters $p = [\mu, \sigma, u_2]$ and ξ are estimated alternatively. Indeed, for each iteration, we estimate the parameters vector p by minimizing the Squared Error (SE) between the empirical cdf and the estimated one, when considering the estimate of ξ of the previous iteration. Thereafter, ξ will be estimated again, as well, by minimizing the SE between the empirical cdf and the estimated one, using, this time, the estimated p vector. Evidently, this procedure starts by fixing initial parameters and ends when a stop condition is satisfied. Small modifications on the parameters vector, the data scale, and the stop condition, are provided when going from the two- to the three-components algorithm, and will be detailed later on. These modifications have no influence on the functional principle of the algorithm, neither on how to study its convergence. In this convergence study, given in the appendix, we prove analytically the existence of a stationary point, then show numerically that the stationary point is attractive and unique. This last part is still an open analytical question.

Let us describe this iterative algorithm. First, we consider an n -sample $X = (X_i)_{1 \leq i \leq n}$ with a G-E-GPD parent distribution. We denote by $x = (x_i)_{1 \leq i \leq n}$ an associated given realization. For the rest of this work, $\tilde{a}^{(0)}$ and $\tilde{a}^{(k)}$ denote the initialization and the estimate of the parameter a at the k^{th} iteration, respectively.

To start its iterative process, the three-components algorithm needs the knowledge of $\tilde{\xi}^{(0)}$, the initialization of ξ . However, the only information we have about ξ is that it is positive, which makes its initialization difficult. For that reason, we start initializing $\tilde{p}^{(0)} = [\tilde{\mu}^{(0)}, \tilde{\sigma}^{(0)}, \tilde{u}_2^{(0)}]$. To do so, we chose $\tilde{\mu}^{(0)}$ as the mode of the data, according to the fact that about 16% of Gaussian observations are below $\mu - \sigma$, we took $\tilde{\sigma}^{(0)} = \tilde{\mu}^{(0)} + q_{16\%}$, where $q_{16\%}$ represents the quantile of order 16% associated to H , and $\tilde{u}_2^{(0)}$ as a quantile of order sufficiently high (above 80%, for this work). Then we use this initialization $\tilde{p}^{(0)}$ to determine $\tilde{\xi}^{(0)}$, minimizing the SE between the hybrid cdf given $p = \tilde{p}^{(0)}$ (fixed), and the empirical cdf H_n associated to the sample $X = (X_i)_{1 \leq i \leq n}$, defined,

for all $t \in \mathbb{R}$, by $H_n(t) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{(X_i \leq t)}$. To do so, we do not evaluate this SE on the realizations x_i only (as there might be only a few observations in the tail), but on a generated sequence of synthetic increasing data $y = (y_j)_{1 \leq j \leq m}$, of size m (m can be different from n), with a logarithmic step, in order to increase the number of points above u_2 . More precisely, for any $1 \leq j \leq m$, y_j is expressed as:

$$y_j = \min_{1 \leq i \leq n}(x_i) + (\max_{1 \leq i \leq n}(x_i) - \min_{1 \leq i \leq n}(x_i)) \log_{10} \left(1 + \frac{9(j-1)}{m-1} \right). \quad (3.1)$$

Notice that the introduction of new points between the observations of X has an impact on H by evaluating it on more points, but not on the step function H_n .

Hence $\tilde{\xi}^{(0)}$ is now determined by solving the following minimization problem using the LM algorithm (see [30, 33]):

$$\tilde{\xi}^{(0)} \leftarrow \underset{\xi > 0}{\operatorname{argmin}} \left\| H(y; \theta \mid \tilde{p}^{(0)}) - H_n(y) \right\|_2^2,$$

where $\theta \mid \tilde{p}^{(0)}$ represents θ for $p = \tilde{p}^{(0)}$.

Note that this initialization step is the first modification we have introduced, compared with the two-components algorithm.

Once $\tilde{\xi}^{(0)}$ is determined, we can, thereafter, proceed iteratively. For all $k \geq 1$, the k^{th} iteration is

split into two main minimization problems, which are solved alternatively, as described hereafter.

Step 1: Determination of $\tilde{p}^{(k)} = [\tilde{\mu}^{(k)}, \tilde{\sigma}^{(k)}, \tilde{u}_2^{(k)}]$, minimizing the SE between the hybrid cdf given $\tilde{\xi}^{(k-1)}$, and the empirical one, as follows:

$$\tilde{p}^{(k)} \leftarrow \underset{\substack{(\mu, \sigma) \in \mathbb{R} \times \mathbb{R}_+^* \\ u_2 \in \mathbb{R}_+}}{\operatorname{argmin}} \left\| H(y; \theta | \tilde{\xi}^{(k-1)}) - H_n(y) \right\|_2^2$$

where $\theta | \tilde{\xi}^{(k-1)}$ denotes θ for $\xi = \tilde{\xi}^{(k-1)}$ (fixed).

This minimization problem is as well numerically resolved using the LM algorithm.

Step 2: Determination of $\tilde{\xi}^{(k)}$, minimizing the SE between the hybrid cdf given $\tilde{p}^{(k)}$, and the empirical one, *i.e.* by solving the following minimization problem via the LM algorithm:

$$\tilde{\xi}^{(k)} \leftarrow \underset{\xi > 0}{\operatorname{argmin}} \left\| H(y; \theta | \tilde{p}^{(k)}) - H_n(y) \right\|_2^2,$$

where $\theta | \tilde{p}^{(k)}$, represents θ for $p = \tilde{p}^{(k)}$ (fixed).

Stop condition: The algorithm iterates until it satisfies the following stop condition:

$$\left(\underbrace{d\left(H(y; \theta^{(k)}), H_n(y)\right) < \epsilon}_{\text{Condition C1}} \quad \text{and} \quad \underbrace{d\left(H(y_{q_\alpha}; \theta^{(k)}), H_n(y_{q_\alpha})\right) < \epsilon}_{\text{Condition C2}} \right) \quad \text{or} \quad \underbrace{k = k_{max}}_{\text{Condition C3}}$$

where $d(a, b)$ denotes the distance between a and b , chosen in this study as the Mean Squared Error (MSE), $\alpha \in [0, 1]$, and y_{q_α} represents the observations above the quantile q_α of order α associated with H .

The second modification with respect to the two-components algorithm initially developed, lies at the stop condition. Indeed, to ensure a reliable fit of data not only for the main behavior but also for the tail, we force the algorithm to stop only when the MSE between the hybrid cdf and the empirical one is small enough ($\epsilon = 10^{-12}$, for this work), using on one hand all data (Condition C1), on the other hand only extreme order statistics above a desired q_α (Condition C2). Otherwise, the algorithm stops when a fixed number k_{max} of iterations ($k_{max} = 10^3$, for this work) is reached (Condition C3).

REMARK 3.1 *Note that this algorithm can be adapted to different hybrid models according to the nature and the number of its components (if larger than 2), without any influence on the convergence study of the adapted algorithm. We point out that in this method, it is important to have at least two-components, among which a GPD to describe the extremes behavior, to be able to determine in an automatic way the threshold above which the GPD is fitted. For simplicity, we focus on the right tail, but it is straightforward to repeat the same procedure to consider the left tail too. The algorithm has been extended in an example of this type, when considering both tails (see [10]). We just follow the same logic: breaking down the problem of parameters estimation into two subproblems and then resolving them alternatively.*

To summarize, let us provide a pseudo-code of our algorithm.

Algorithm 1 Iterative and unsupervised algorithm for the G-E-GPD parameters estimation

1: Initialization of $\tilde{p}^{(0)}$, α , $\epsilon > 0$, and k_{max} , then initialization of $\tilde{\xi}^{(0)}$:

$$\tilde{\xi}^{(0)} \leftarrow \underset{\xi > 0}{\operatorname{argmin}} \left\| H(y; \theta \mid \tilde{p}^{(0)}) - H_n(y) \right\|_2^2,$$

where H_n is the empirical cdf of X . We note that this distance is computed on the points $y = (y_j)_{1 \leq j \leq m}$ defined in (3.1).

2: Iterative process:

• $k \leftarrow 1$

Step 1 - Estimation of $\tilde{p}^{(k)}$:

$$\tilde{p}^{(k)} \leftarrow \underset{\substack{(\mu, \sigma) \in \mathbb{R} \times \mathbb{R}_+^* \\ u_2 \in \mathbb{R}_+}}{\operatorname{argmin}} \left\| H(y; \theta \mid \tilde{\xi}^{(k-1)}) - H_n(y) \right\|_2^2$$

Step 2 - Estimation of $\tilde{\xi}^{(k)}$:

$$\tilde{\xi}^{(k)} \leftarrow \underset{\xi > 0}{\operatorname{argmin}} \left\| H(y; \theta \mid \tilde{p}^{(k)}) - H_n(y) \right\|_2^2,$$

• $k \leftarrow k + 1$

until $(d(H(y; \theta^{(k)}), H_n(y)) < \epsilon$ and $d(H(y_{q_\alpha}; \theta^{(k)}), H_n(y_{q_\alpha})) < \epsilon)$ or $(k = k_{max})$.

3: Return $\theta^{(k)} = [\tilde{\mu}^{(k)}, \tilde{\sigma}^{(k)}, \tilde{u}_2^{(k)}, \tilde{\xi}^{(k)}]$.

4 Simulation results and discussion

To study the performance of the algorithm to self-calibrate the G-E-GPD model, we build on Monte-Carlo simulations. To do so, we proceed in 4 steps:

Step 1: We consider N training sets $\{x^q = (x_p^q)_{1 \leq p \leq n}\}_{1 \leq q \leq N}$, of length n and N test sets $\{y^q = (y_p^q)_{1 \leq p \leq l}\}_{1 \leq q \leq N}$, of length l , with a G-E-GPD parent distribution admitting a fixed parameters vector θ .

Step 2: On each training set x^q , $1 \leq q \leq N$, we estimate θ , say $\tilde{\theta}^q = [\tilde{\mu}^q, \tilde{\sigma}^q, \tilde{u}_2^q, \tilde{\xi}^q]$, using the algorithm given in the previous section. We denote by \tilde{a}^q the estimation of the parameter a relative to the q^{th} training set.

Step 3: We compute the empirical mean and variance of estimates of each parameter over the N training sets. For any parameter a , we denote by \tilde{a} and \tilde{S}_a its empirical mean and variance, respectively, defined as: $\tilde{a} = \frac{1}{N} \sum_{q=1}^N \tilde{a}^q$ and $\tilde{S}_a = \frac{1}{N-1} \sum_{q=1}^N (\tilde{a}^q - \tilde{a})^2$. We can check the relevance of \tilde{a} using two criterions:

1. The MSE expressed for any parameter a as: $\text{MSE}_a = \frac{1}{N} \sum_{q=1}^N (\tilde{a}^q - a)^2$. A small value of

MSE highlights the reliability of parameters estimation using the presented algorithm.

2. Test on the mean (with unknown variance) : $\left| \begin{array}{ll} H_0 & : \tilde{a} = a \\ H_1 & : \tilde{a} \neq a \end{array} \right.$

For $N > 30$, we can use a normal test (instead of a t -test) of size α , with a rejection region of H_0 at level α described by $(|T_{\tilde{a}}| > \Phi^{-1}(1 - \alpha))$, where the statistics $T_{\tilde{a}}$ is given

by $T_{\tilde{a}} = \sqrt{\frac{N}{\tilde{S}_N^a}} (\tilde{a} - a)$, and $\Phi^{-1}(1 - \alpha)$ is the quantile of order $1 - \alpha$ of the standard normal distribution.

Step 4: We compare the hybrid pdf h given θ with the pdf \tilde{h} estimated on each test set y^q , given $\tilde{\theta}^q$. To do so, we compute the average of the log-likelihood ratio \mathcal{D} of $h(y^q; \tilde{\theta}^q)$ by $h(y^q; \theta)$, over the N simulations:

$$\mathcal{D} = \frac{1}{Nl} \sum_{q=1}^N \sum_{p=1}^l \log \left(\frac{h(y_p^q; \theta)}{\tilde{h}(y_p^q; \tilde{\theta}^q)} \right). \quad (4.1)$$

It is obvious that the smallest the value of \mathcal{D} is, the most trustworthy is the algorithm.

We present in Table 1 the results obtained when taking $\theta = [2, 1, 5, 0.5]$, $N = 100$, $l = n$ and $\alpha = 5\%$. Different values of n have been considered to study its impact on the parameters estimation. The reliability of the three-components algorithm, in terms of goodness-of-fit, is pointed out through the two criterions described above, as well as *via* the average of the log-likelihood ratio. First, for each estimated parameter, we notice a small MSE whenever the data size is large enough, with a variance of order $1/n$ (except for u_2 , where it is much larger). This order being larger than standard ones ($1/n^2$, as *e.g.* for the the Hill and QQ estimators), we resort to a statistical test as an additional criterion. For the N training sets, we compute the test statistics denoted $T_{\tilde{a},N}$ and the corresponding p -value $p_{T_{\tilde{a},N}} = 2(1 - \Phi(|T_{\tilde{a},N}|))$, with respect to the parameter a . For any $n \in \{10^3, 10^4, 10^5\}$ and for any parameter $a \in \{\mu, \sigma, u_2, \sigma\}$, we obtain $|T_{\tilde{a},N}| < \Phi^{-1}(0.95) = 1.6448$, and $p_{T_{\tilde{a},N}} > \frac{\alpha}{2} = 2.5\%$ (it even remains greater than 50%), which reveals a high acceptance (95% level) of H_0 ($\tilde{a} = a$) *i.e.* a very high level of similarity between the values obtained via the algorithm and the fixed ones (even for $\alpha = 1\%$). Finally, the accuracy of the parameters estimation is also highlighted via the average log-likelihood ratio. For the three different values of n , this average is lower than 10^{-2} , involving a good self-calibration of the G-E-GPD hybrid model.

A remaining question, which might be the object of another paper, is the study of the convergence rate of this algorithm. Here, to have an idea of how fast it works, we indicate in Table 1 the average execution time and the average iterations number (the floor function) over the N simulations. As shown in this table, they both increase with the data size, as expected. We notice that the average execution time is still small, even for $n = 10^5$, indicating a fast convergence of the algorithm. It could be even reduced by converting our programs from the R programming language to the C++ one.

Besides the reliable estimation of the parameters, we show in Table 2, *via* the MSE, that our algorithm enhances the GPD parameters estimation when compared with the Maximum Likelihood (ML) method and the Probability Weighted Moments (PWM) (see [22]). We mention that the threshold we select for the comparison is the one obtained by the algorithm.

Table 1: Monte-Carlo simulations results for $\theta = [2, 1, 5, 0.5]$, $N = 100$, $\alpha = 5\%$, and $l = n \in \{10^3, 10^4, 10^5\}$.

			$n = 10^3$	$n = 10^4$	$n = 10^5$
Parameters	$\mu = 2$	$\widetilde{\mu}$	1.9981	1.9994	1.9994
		\widetilde{S}_N^μ	$6.8722 \cdot 10^{-3}$	$8.9157 \cdot 10^{-4}$	$7.6821 \cdot 10^{-5}$
		MSE_μ	$6.8071 \cdot 10^{-3}$	$8.8295 \cdot 10^{-4}$	$7.6143 \cdot 10^{-5}$
		$T_{\widetilde{\mu}, N}$	-0.2285	-0.1829	-0.3434
		$pT_{\widetilde{\mu}, N}$	0.8192	0.8548	0.7312
	$\sigma = 1$	$\widetilde{\sigma}$	1.0013	1.0007	0.9999
		\widetilde{S}_N^σ	$4.739 \cdot 10^{-3}$	$4.8821 \cdot 10^{-4}$	$5.231 \cdot 10^{-5}$
		MSE_σ	$4.6934 \cdot 10^{-3}$	$4.8386 \cdot 10^{-4}$	$5.1791 \cdot 10^{-5}$
		$T_{\widetilde{\sigma}, N}$	0.1926	0.3307	-0.083
		$pT_{\widetilde{\sigma}, N}$	0.8472	0.7408	0.9338
	$u_2 = 5$	$\widetilde{u_2}$	4.9904	4.9896	4.9964
		$\widetilde{S}_N^{u_2}$	$5.4904 \cdot 10^{-1}$	$4.9062 \cdot 10^{-2}$	$3.4532 \cdot 10^{-3}$
		MSE_{u_2}	$5.4364 \cdot 10^{-1}$	$4.8678 \cdot 10^{-2}$	$3.4311 \cdot 10^{-3}$
		$T_{\widetilde{u_2}, N}$	-0.1288	-0.4669	-0.5996
		$pT_{\widetilde{u_2}, N}$	0.8975	0.6405	0.5487
	$\xi = 0.5$	$\widetilde{\xi}$	0.4975	0.5005	5.0018
		\widetilde{S}_N^ξ	$1.6698 \cdot 10^{-3}$	$1.4635 \cdot 10^{-4}$	$1.113 \cdot 10^{-5}$
		MSE_ξ	$1.6594 \cdot 10^{-3}$	$1.452 \cdot 10^{-4}$	$1.1053 \cdot 10^{-5}$
		$T_{\widetilde{\xi}, N}$	-0.6102	0.4592	0.5478
		$pT_{\widetilde{\xi}, N}$	0.5416	0.646	0.5837
Average execution time (seconds)			3.831	13.2048	245.6802
Average iterations number			45	48	50
\mathcal{D}			$2.9858 \cdot 10^{-3}$	$2.6957 \cdot 10^{-4}$	$2.9898 \cdot 10^{-5}$

Table 2: GPD parameters estimation, using Algorithm 1, for the same example ($\theta = [2, 1, 5, 0.5]$, $N = 100$, and $\beta = 2.5$) and for $n = 10^5$: comparison with the ML and the PWM methods.

		Algorithm 1	ML	PWM
$\xi = 0.5$	$\tilde{\xi}$	5.0018	0.5003	0.5025
	MSE_ξ	$1.1053 \cdot 10^{-5}$	$1.7312 \cdot 10^{-4}$	$5.0868 \cdot 10^{-4}$
$\beta = 2.5$	$\tilde{\beta}$	2.499	2.5003	2.4968
	MSE_β	$4.0818 \cdot 10^{-4}$	$1.3788 \cdot 10^{-3}$	$2.3523 \cdot 10^{-3}$

5 Application of the self-calibrating method on real data

Once the performance of the algorithm is validated on generated data, we apply it on real data, considering two different domains: neuroscience and finance. Those data are essentially symmetric around the mean. However, in the case of skewed data, as in insurance claims, the normal component would be replaced by a lognormal one, to account for this skewness (as e.g. in [25]). To underline the unsupervised aspect of the self-calibrating method, for each application we compare, in terms of goodness-of-fit of extremes, the results obtained with this method to those provided by three standard EVT approaches: the graphical Mean Excess Plot (MEP) (see [15]), Hill (see [21]) and QQ (see [28]) methods.

5.1 Neuroscience: neural data

Here we consider the data corresponding to twenty seconds, equivalent to $n = 3 \cdot 10^5$ observations, of real extracellular recording of neurons activities, available in [42] and measured on the antennal lobe of an adult locust (see [43]). We represent one second of the considered data in Figure 1. The information to be extracted from these data (spikes or action potentials; see [34]) lies on the extreme behaviors (left and right) of the data.




Figure 1: One second of neural data, extracellularly recorded.

Since the neural data can be considered as symmetric, it is sufficient to evaluate the right side of the distribution with respect to its mode. We compare in the following the obtained results of neural data fitting (only the right side of the data) using our self-calibrating method, the MEP, the Hill and the QQ ones.

- *Application of our self-calibrating method*

Applying the algorithm of this method on the neural data set to model its right side, we obtain the following estimate of θ : $\hat{\theta} = [-0.0681, 0.6297, 0.5398, 1.0301]$. In Figure 2 (1st row), we can see well, on a log-scale, the good fit of the estimated hybrid cdf compared to the empirical one (see plot (a)), but also that of the right tail distribution (see plot (b)). We observe here that the exponential distribution is not needed for a good modeling of the data. Indeed, the two junction points overlap: \tilde{u}_1 and \tilde{u}_2 , the estimates of u_1 and u_2 , respectively,

are very close to each other (with a distance equal to $4.3268 \cdot 10^{-5}$) (see plot (a)). It is what has already been observed in [12].

- *Application of the MEP method*

We draw the Mean Excess Plot (MEP) (see plot (d)) to manually determine the threshold above which data are GPD distributed (that we refer as the GPD threshold). We look from which threshold (high enough to match the theory, but not too high to have enough observations) the MEP behaves linearly. Then, we estimate the corresponding GPD parameters using, for instance, the PMW method (see [22]). Several values of the threshold have been selected. We choose the one offering the smallest MSE between the empirical tail distribution and the estimated GPD (see the zoomed part of plot (c), where the linear behavior of the MEP is pointed out). The reliability of this graphical method, in terms of goodness-of-fit of extremes above the selected threshold, is illustrated in plot (d).

- *Application of the Hill method*


In a similar way, we determine the GPD threshold graphically from the Hill plot (see plot (e)), representing the Hill estimator of the GPD tail index as a function of the number of the upper order statistics. After several tests, we select the number of upper order statistics above which we observe a stability of the Hill plot (see plot (e)); the associated threshold minimizes the MSE between the estimated and the empirical tail cdf. We note that once the number of upper order statistics is selected, the associated threshold and tail index are determined, and the scale parameter is estimated as the product of this threshold by the tail index. We draw in plot (f) the empirical tail distribution and the estimated GPD.

- *Application of the QQ-method*


We provide not only the Hill estimator but also the QQ one, as it is easier to detect the threshold with the QQ method than with the Hill plot, which may sometimes present an erratic behavior difficult to interpret, as can be observed in the corresponding zoomed plots. As well as for the Hill method, we select graphically the number of upper order statistics above which the plot of the QQ-estimator of the tail index (see plot (g)) behaves this time linearly. We then select the one such that the MSE between the empirical tail distribution and the estimated GPD is minimum (see the zoomed part of plot (g)). As for the Hill method, the scale parameter is estimated as the product of the threshold by the tail index. Plot (h) illustrates the reliability of the obtained tail fit.

- *Comparison of the results obtained via the various methods*


In Table 3, we present the results we obtained with the self-calibrating method, the MEP, Hill and QQ methods. Since the three graphical approaches fit only the tail distribution, the comparison of the methods will focus on the goodness-of-fit of the GPD component. As observed in this table, the MSE between the estimated cdf and the empirical one, using only data above the selected threshold is small enough for the four methods to ensure a reliable modeling of extremes. The GPD threshold and the estimated tail index are of the same order of magnitude for all methods; it confirms that our algorithm works in the right direction. We can also notice the good performance of these methods through Figure 3, where we plot the empirical quantile function and the estimated ones using the self-calibrating method and the various graphical ones. However, the advantage of our method is that it is unsupervised,




(a) Right side of the empirical cdf vs the hybrid cdf obtained with the self-calibrating method




(b) Empirical tail distribution vs GPD - Self calibrating method




(c) Mean Excess Plot




(d) Empirical tail distribution vs GPD - MEP & PWM method




(e) Hill plot



(f) Empirical tail distribution vs GPD - Hill method



(g) QQ-estimator plot



(h) Empirical tail distribution vs GPD - QQ method

Figure 2: Extremes modeling using different methods. For each plot, the blue continuous curve is empirical (even for zoomed curves), while the red dashed curve and vertical line represent the estimated GPD and threshold, respectively, using the associated method.

i.e. it does not need the intervention of the user to select the threshold manually. Moreover it provides a good fit between the hybrid cdf estimated on the entire data sample and the empirical cdf, with a MSE of order 10^{-5} .

Table 3: Comparison between the self-calibrating method and the three graphical methods: MEP, Hill and QQ ones. N_{u_2} represents the number of observations above u_2 . The distance gives the MSE between the empirical (tail or full respectively) distribution and the estimated one from a given model (GPD or hybrid G-E-GPD respectively). The neural data sample size is $n = 3 \cdot 10^5$.

Model	tail index (ξ)	threshold (u_2)	N_{u_2}	distance (tail distr.)	distance (full distr.)
GPD	MEP (PWM): 0.3326	$1.0855 = q_{93.64\%}$	19260	$4.0663 \cdot 10^{-6}$	
GPD	Hill-estimator: 0.599	$1.0855 = q_{93.64\%}$	19260	$2.0797 \cdot 10^{-6}$	
GPD	QQ-estimator: 0.5104	$1.0671 = q_{93.47\%}$	19871	$1.2685 \cdot 10^{-5}$	
G-E-GPD	Self-calibrating method: 0.5398	$1.0301 = q_{92.9\%}$	21272	$7.7903 \cdot 10^{-6}$	$9.3168 \cdot 10^{-5}$




Figure 3: Neural data: Comparison between the empirical quantile function and the estimated ones, the self-calibrating methods and the graphical methods.

5.2 Finance: S&P500 absolute log-returns

The second application considered in this work concerns the S&P500 absolute log-returns from January 2, 1987 to February 29, 2016, corresponding to $n = 7349$ observations, available in the tseries package (see [52]) of the R programming language (see [44]). We check whether these data, represented in Figure 4, exhibit a heavy tail or not. Note that we could have looked at each tail of

the log-returns (it is known that they have a slight different tail index from one to the other) but decided instead to look at their absolute value.




Figure 4: *S&P500 absolute daily log-returns from January 2, 1987 to February 29, 2016*

As for the neural data, we apply our self-calibrating method, and the graphical ones (MEP and QQ) for comparison. Note that we display only the plots associated with the QQ-method and not the Hill one, since the QQ-threshold is easier to detect, as already commented. Nevertheless, we provide the numerical results for both Hill and QQ methods, using the \sqrt{n} upper order statistics to compute the Hill estimator, since this selected threshold has been empirically shown to be relevant for financial data in [3].

The results are illustrated in Figure 5. In plot (a), we draw (on a log-scale) the empirical cdf and the hybrid cdf obtained via our self-calibrating method, where the two vertical dashed lines represent the two junction points of the hybrid model. The corresponding right tail fit is given in plot (b).

The MEP and the selected threshold are given in plot (c) (the zoomed part shows the MEP linear behavior above the selected threshold), while the corresponding extremes fit is given in plot (d). The GPD parameters being estimated by the PWM method.


Finally, the QQ-estimator plot and the selected number of upper order statistics are represented in plot (e), with a zoom illustrating the linear behavior of the QQ-estimator plot above the selected number of upper order statistics. The corresponding extremes fit is shown in plot (f).

Table 4: *Comparison between the self-calibrating method and the three graphical methods: MEP, Hill and QQ ones. The S&P 500 absolute log-returns data sample size is $n = 7349$.*


Model	tail index (ξ)	threshold (u_2)	N_u	distance (tail distr.)	distance (full distr.)
GPD	MEP: 0.3025	$0.0282 = q_{97.21\%}$	206	$1.7811 \cdot 10^{-7}$	
GPD	Hill-estimator: 0.3094	$0.0382 = q_{98.85\%}$	85	$4.4953 \cdot 10^{-8}$	
GPD	QQ-estimator: 0.3288	$0.0323 = q_{98.14\%}$	137	$6.0505 \cdot 10^{-8}$	
G-E-GPD	Self-calibrating method: 0.3332	$0.0289 = q_{97.49\%}$	184	$1.9553 \cdot 10^{-7}$	$1.0635 \cdot 10^{-5}$

The numerical results obtained for the threshold and tail index, as well as for the MSE between the empirical tail distribution and the estimated GPD using the four methods respectively, are reported in Table 4. We can notice that all methods offer a good fit of the tail distribution, with a slightly overestimation for the G-E-GPD and QQ methods compared with the MEP and Hill


ones. In Figure 6, we also give a comparison of the estimated quantile function using the G-E-GPD method and the graphical (MEP, Hill and QQ) ones.




(a) Empirical cdf vs the hybrid cdf obtained with the self-calibrating method




(b) Empirical tail distribution vs GPD - Self calibrating method




(c) Mean Excess Plot



(d) Empirical tail distribution vs GPD - MEP & PWM method



(e) QQ-estimator plot



(f) Empirical tail distribution vs GPD - QQ method

Figure 5: Extremes modeling using different methods. For each plot, the blue continuous curve is empirical (even for zoomed curves), while the red dashed curve and vertical line represent the estimated GPD and threshold, respectively, using the associated method.




Figure 6: S&P 500 absolute log-returns data: Comparison between the empirical quantile function and the estimated ones, the self-calibrating methods and the graphical methods.

In Table 4 and Figure 6, we observe once again similar results for the various methods. It confirms the good performance of the self-calibrating method to estimate the tail distribution. As already said, this latter method also provides a good modeling for the entire cdf.

Conclusion

In this paper, we propose a self-calibrating method to model heavy tailed data that may be non-homogeneous and multi-components. We develop it introducing a general non-degenerate hybrid \mathcal{C}^1 distribution for heavy tailed data modeling, which links a normal distribution to a GPD via an exponential distribution that bridges the gap between mean and asymptotic behaviors. The three distributions are connected to each other at junction points estimated by an iterative algorithm, as are the other parameters of the model. The convergence of the algorithm is studied analytically for one part and numerically for the other. The performance of the method is studied on simulated data. Based on those results, we observe that the proposed unsupervised algorithm offers a judicious fit of the asymmetric right heavy tailed data with an accurate determination of the tail threshold indicating the presence of extremes, as well as of the parameters of the GPD that fits the extremes over this threshold. Several applications of the method have been done on real data, in particular on insurance data. We give two of them on data coming from very different fields, neural data and financial ones (S&P500). A comparison follows with other existing methods.

Note that this method has been developed when considering asymmetric right heavy tailed data; it can of course be applied in the same way when having the asymmetry on the left side, or when having a heavy tail on each side (without requiring a symmetry).

This method has many advantages and should become of great use in practice. The main advantage is to be unsupervised, avoiding the somehow arbitrary resort, when fitting the tail, to standard graphical methods (e.g. MEP, Hill, QQ methods) in EVT. A second advantage is to fit with the same iterative algorithm the full distribution of observed heavy tailed data, of any type whenever smooth enough (\mathcal{C}^1 -distribution), providing an accurate estimation of the parameters for the mean and extreme behaviors. It certainly answers a big concern encountered by practitioners. Moreover the method is quite general: besides the GPD needed when fitting the heavy tail, the other components might be chosen differently, not using limit behavior (CLT) but distributions chosen specifically for the data that are worked out (as e.g. lognormal for insurance claims). It would not change at all the structure of the algorithm.

It should be emphasized that determining in a unsupervised way the threshold over which we have extremes, requires to have information before the threshold. We suggest here an approach that avoids traditional graphical methods when fitting the entire distribution. Further investigation will follow in order to make this method also available as a pure EVT tool (i.e. to fit the tail only). It means to determine the minimum information required to determine the neighbor distribution of the GPD to have a robust estimation for the tail threshold and the GPD parameters estimation.

Moreover, we plan also to tackle the analytical study of the convergence rate of the algorithm as a function of sample size.

Finally, a R package should appear soon online. Meantime, the R codes are available upon request.

Acknowledgments

The first two authors acknowledge the support from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no 318984 - RARE.

References

- [1] BEIRLANT, J., GOEGBEUR, Y., SEGERS, J., TEUGELS, J., DE WAAL, D., AND FERRO, C. *Statistics of Extremes: Theory and Applications*. Wiley Series in Probability and Statistics. Wiley, 2006.
- [2] BERNARDO, J. M., AND SMITH, A. F. *Bayesian Theory*, vol. 405. John Wiley & Sons, 2009.
- [3] BLUM, P., AND DACOROGNA, M. Extreme forex moves. *Risk-London-Risk Magazine Limited* 16, 2 (2003), 63–66.
- [4] BROADWATER, J. B., AND CHELLAPPA, R. Adaptive threshold estimation via extreme value theory. *IEEE Transactions on Signal Processing* 58, 2 (2010), 490–500.
- [5] CARREAU, J., AND BENGIO, Y. A hybrid pareto model for asymmetric fat-tailed data: the univariate case. *Extremes* 12, 1 (2009), 53–76.

- [6] DACOROGNA, M. M., AND KRATZ, M. Living in a stochastic world and managing complex risks. *Available at SSRN 2668468* (2015).
- [7] DAVISON, A. C., AND SMITH, R. L. Models for exceedances over high thresholds. *Journal of the Royal Statistical Society. Series B (Methodological)* (1990), 393–442.
- [8] DE HAAN, L., AND FERREIRA, A. *Extreme Value Theory: An Introduction*. Springer Science & Business Media, 2006.
- [9] DEBBABI, N. *Approche algébrique et théorie des valeurs extrêmes pour la détection de ruptures: Application aux signaux biomédicaux*. PhD thesis, Université de Reims Champagne Ardenne, France, 2015.
- [10] DEBBABI, N., EL ASMI, S., AND MBOUP, M. Distribution hybride pour la modélisation de données à deux queues lourdes: Application sur les données neuronales. In *Groupe d’Etudes du Traitement du Signal et des Images, GRETSI* (2015).
- [11] DEBBABI, N., AND KRATZ, M. A new unsupervised threshold determination for hybrid models. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2014), pp. 3440–3444.
- [12] DEBBABI, N., KRATZ, M., AND EL ASMI, S. A hybrid distribution for highly right skewed data. In *2015 International Conference on Communications and Networking (ComNet)* (2015).
- [13] DEBBABI, N., KRATZ, M., AND MBOUP, M. Unsupervised threshold determination for hybrid models. Tech. rep., University of Reims, june 2013.
- [14] DEBBABI, N., KRATZ, M., MBOUP, M., AND ASMI, S. E. Combining algebraic approach with extreme value theory for spike detection. In *Signal Processing Conference (EUSIPCO), 2012 Proceedings of the 20th European* (2012), pp. 1836–1840.
- [15] EMBRECHTS, P., KLÜPPELBERG, C., AND MIKOSCH, T. *Modelling Extremal Events for Insurance and Finance*. Springer-Verlag, 1997.
- [16] FELDMAN, R., AND TAQQU, M. *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*. Springer Science & Business Media, 1998.
- [17] FINE, T., AND LEVIN, M. Extreme-value theory applied to false-alarm probabilities (corresp.). *IRE Transactions on Information Theory* 8, 3 (1962), 259–260.
- [18] GENÇAY, R., DACOROGNA, M., MULLER, U. A., PICTET, O., AND OLSEN, R. *An Introduction to High-frequency Finance*. Academic Press, 2001.
- [19] GUILLOU, A., KRATZ, M., AND STRAT, Y. L. An extreme value theory approach for the early detection of time clusters. a simulation-based assessment and an illustration to the surveillance of salmonella. *Statistics in medicine* 33, 28 (2014), 5015–5027.
- [20] GUMBEL, E. *Statistics of extremes*. Columbia University Press, 1958.
- [21] HILL, B. M. A simple general approach to inference about the tail of a distribution. *The Annals of Statistics* 3 (1975), 1163–1174.
- [22] HOSKING, J., AND WALLIS, J. Parameter and quantile estimation for the generalized pareto distribution funcrion. *Technometric* 29 (3) (1987), 339–349.
- [23] JOOSEUK, K., AND CLAYTON, D. Robust kernel density estimation. *Journal of Machine*

- Learning Research* 13 (2012), 2529–2565.
- [24] KAKUTANI, S. A generalization of brouwer’s fixed point theorem. *Duke Mathematical Journal* 8, 3 (1941), 457–459.
 - [25] KNECHT, M., AND KÜTTEL, S. The Czeledin distribution function. In *XXXIV ASTIN Colloquium, Berlin* (2003).
 - [26] KOLLU, R., RAYAPUDI, S. R., NARASIMHAM, S., AND PAKKURTHI, K. M. Mixture probability distribution functions to model wind speed distributions. *International Journal of Energy and Environmental Engineering* 3, 1 (2012), 1–10.
 - [27] KRATZ, M. Normex, a new method for evaluating the distribution of aggregated heavy tailed risks. *Extremes*, 4 (2014). Special issue on *Extremes in Finance* (Guess Ed. P. Embrechts): 661–691.
 - [28] KRATZ, M., AND RESNICK, S. I. The qq-estimator and heavy tails. *Communications in Statistics. Stochastic Models* 12, 4 (1996), 699–724.
 - [29] LEADBETTER, M. R., LINDGREN, G., AND ROOTZÉN, H. *Extremes and Related Properties of Random Sequences and Processes*. 3Island Press, 1983.
 - [30] LEVENBERG, K. A method for the solution of certain nonlinear problems in least squares quart. *Applied Math* 2 (1944), 164–168.
 - [31] LI, C., SINGH, V. P., AND MISHRA, A. K. Simulation of the entire range of daily precipitation using a hybrid probability distribution. *Water resources research* 48 (2012), 1–17.
 - [32] MANDAVA, A. K., SHAHRAM, L., AND REGENTOVA, E. E. Reliability assessment of microarray data using fuzzy classification methods: A comparative study. *Advances in Computing and Communications* (2011), 351–360.
 - [33] MARQUARDT, D. W. An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal on Applied Mathematics* 11 (1963), 431–441.
 - [34] MBOUP, M. Neural spike detection and localisation via volterra filtering. In *2012 IEEE International Workshop on Machine Learning for Signal Processing (MLSP)* (2012).
 - [35] MCNEIL, A., FREY, R., AND EMBRECHTS, P. *Quantitative Risk Management: Concepts, Techniques, and Tools*. Princeton Series in Finance. Princeton University Press, 2005.
 - [36] MILSTEIN, L., SCHILLING, D., AND WOLF, J. Robust detection using extreme-value theory. *IEEE Transactions on Information Theory* 15, 3 (1969), 370–375.
 - [37] NADARAJAH, S., AND BAKAR, S. New composite models for the danish fire insurance data. *Scandinavian Actuarial Journal* 2014, 2 (2014), 180–187.
 - [38] NICULESCU, C. P., AND ROVENȚA, I. Schauder fixed point theorem in spaces with global nonpositive curvature. *Fixed Point Theory and Applications* 1 (2009), 1–8.
 - [39] ORBANZ, P., AND TEH, Y. W. *Bayesian Nonparametric Models*. Springer US, 2010, pp. 81–89.
 - [40] PARZEN, E. On estimation of a probability density function and mode. *The annals of mathematical statistics* 33, 3 (1962), 1065–1076.
 - [41] PICKANDS, J. Statistical inference using extreme order statistics. *Annals of Statistics* 3 (1975),

- 119–131.
- [42] POUZAT, C. Tetrode recording from the antennal lobe of a locust (*Schistocerca americana*), Jan. 2015. <http://dx.doi.org/10.5281/zenodo.14607>.
 - [43] POUZAT, C., MAZOR, O., AND LAURENT, G. Using noise signature to optimize spike-sorting and to assess neuronal classification quality. *Journal of neuroscience methods* 122, 1 (2002), 43–57.
 - [44] R CORE TEAM. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2014.
 - [45] RAMSAY, J. O., HOOKER, G., CAMPBELL, D., AND CAO, J. Parameter estimation for differential equations: a generalized smoothing approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69, 5 (2007), 741–796.
 - [46] RANGASWAMY, M., MICHELS, J. H., AND HIMED, B. Statistical analysis of the non-homogeneity detector for {STAP} applications. *Digital Signal Processing* 14, 3 (2004), 253–267.
 - [47] RESNICK, S. I. *Extreme Values, Regular Variation, and Point Processes*. Springer-Verlag, 1987.
 - [48] RESNICK, S. I. *Heavy-tail Phenomena: Probabilistic and Statistical Modeling*. Springer Science & Business Media, 2007.
 - [49] ROSSI, F., FIORENTINO, M., AND VERSACE, P. Two-component extreme value distribution for flood frequency analysis. *Water Resources Research* 20, 7 (1984), 847–856.
 - [50] RUDIN, W. *Principles of Mathematical Analysis*, vol. 3. McGraw-Hill New York, 1964.
 - [51] SERMPEZIS, P., AND SPYROPOULOS, T. Modelling and analysis of communication traffic heterogeneity in opportunistic networks. *IEEE Transactions on Mobile Computing* 14, 11 (2015), 2316–2331.
 - [52] TRAPLETTI, A., AND HORNIK, K. *tseries: Time Series Analysis and Computational Finance*, 2016. R package version 0.10-35.
 - [53] WALKER, S. G., DAMIEN, P., LAUD, P. W., AND SMITH, A. F. M. Bayesian nonparametric inference for random distributions and related functions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 61, 3 (1999), 485–527.
 - [54] WASSERMAN, L. *All of Statistics: A Concise Course In Statistical Inference*. Springer Science & Business Media, 2013.
 - [55] WERNER, D. A proof of the Markov-Kakutani fixed point theorem via the Hahn-Banach theorem. *Extracta Math* 8 (1993), 37–38.
 - [56] WORDEN, K., ALLEN, D. W., SOHN, H., AND FARRAR, C. R. Damage detection in mechanical structures using extreme value statistics. In *SPIE’s 9th Annual International Symposium on Smart Structures and Materials* (2002), International Society for Optics and Photonics, pp. 289–299.
 - [57] ZEYADA, F., MOUBARAK, M., AND SOLIMAN, A. A generalization of generalized Banach fixed-point theorem in a weak left small self-distance space. *Miskolc Mathematical Notes* 8, 99 (2007), 105.

APPENDIX: STUDY OF THE ALGORITHM CONVERGENCE

As already commented, the algorithm convergence does not depend on the number (≥ 2) of components. Therefore, we develop its analysis when considering two-components (*i.e.* $u_1 = u_2$; no exponential component), with a unique weight associated to each one. In the following, we denote by u the junction point connecting the Gaussian distribution to the GPD. We mention that for this two-components hybrid model, named G-GPD, the constraint $\beta = \xi u$ can be relaxed. The parameters vector of the G-GPD model is $\theta = [\mu, \sigma, u]$. The two-components algorithm (see [11]) estimates the parameters $p = [\mu, \sigma]$ and u alternatively. Let us give its pseudo-code for more clarity.

Algorithm 2 Iterative and unsupervised algorithm for the G-GPD parameters estimation

- 1: Initialization of $\tilde{u}^{(0)}$, $\epsilon > 0$, and k_{max} .
- 2: Determination of the empirical cdf H_n associated with our sample $x = (x_i)_{1 \leq i \leq n}$.
- 3: Iterative process:

- $k \leftarrow 1$

Step 1 - Estimation of $\tilde{p}^{(k)} = [\tilde{\mu}^{(k)}, \tilde{\sigma}^{(k)}]$:

$$\tilde{p}^{(k)} \leftarrow \underset{p \in \mathcal{D}_p}{\operatorname{argmin}} \left\| H(x; \theta \mid \tilde{u}^{(k-1)}) - H_n(x) \right\|_2^2,$$

where $\theta \mid \tilde{u}^{(k-1)}$ represents θ for a fixed $u = \tilde{u}^{(k-1)}$, and \mathcal{D}_p is the domain of p for the realization x .

Step 2 - Estimation of $\tilde{u}^{(k)}$:

$$\tilde{u}^{(k)} \leftarrow \underset{u \in \mathcal{D}_u}{\operatorname{argmin}} \left\| H(x; \theta \mid \tilde{p}^{(k)}) - H_n(x) \right\|_2^2,$$

where $\theta \mid \tilde{p}^{(k)}$ means θ for $p = \tilde{p}^{(k)}$ (fixed), and \mathcal{D}_u is the domain of u according to x .

- $k \leftarrow k + 1$
until $(|\tilde{u}^{(k)} - \tilde{u}^{(k-1)}| < \epsilon)$ or $(k = k_{max})$.

- 4: Return $\theta^{(k)} = [\tilde{u}^{(k)}, \tilde{\mu}^{(k)}, \tilde{\sigma}^{(k)}]$.
-

The convergence study is in two main steps. The first one gives the analytical proof of the existence of stationary points. Indeed the algorithm, which consists of a sequence of minimization, does not rely on the optimization of a cost function by seeking a trajectory to reach an extremum of an error surface. As a consequence, the existence of a stationary point, even the convergence towards such one, is not guaranteed, and has to be proved (see Appendix A). The second step consists in checking that the algorithm converges to a unique stationary point. It is done numerically, performing various simulations changing each time the initialization (see Appendix B). We observe that, whatever the initialization, the algorithm converges to the same stationary point. The analytical proof of this second step is still an open problem.

A Existence of stationary points

We start this section by presenting the theoretical framework in which the existence of stationary points has been proved. For a given realization $x = (x_i)_{1 \leq i \leq n}$ and given parameters $\kappa, \tau \in \{p, u\}$ with $\kappa \neq \tau$, we consider the function:

$$\begin{aligned}\varphi_\kappa : \mathcal{D}_\tau &\rightarrow \mathcal{D}_\kappa \\ \tau &\mapsto \varphi_\kappa(\tau; x) = \underset{\kappa \in \mathcal{D}_\kappa}{\operatorname{argmin}} \mathcal{S}_\tau(\kappa; x),\end{aligned}$$

where for $\theta \mid \tau$ denoting, as previously, θ for a given τ , \mathcal{S}_τ is defined by:

$$\mathcal{S}_\tau : \mathcal{D}_\kappa \rightarrow \mathbb{R} \\ \kappa \mapsto \mathcal{S}_\tau(\kappa; x) = \sum_{i=1}^n (H(x_i; \theta \mid \tau) - H_n(x_i))^2$$

To check that φ_κ is a map, it is enough to show that \mathcal{S}_τ admits a unique minimum, for any $\tau \in \{p, u\}$, with $p = [\mu, \sigma] \in \mathbb{R} \times \mathbb{R}_+^*$ and $u \in \mathbb{R}_+$. Since the expression of the hybrid cdf H with respect to (w.r.t.) τ (see (A.1)) is rather complicated, we proceed *via* simulations. Fixing τ , for instance $\tau = p$ (it would be the same for $\tau = u$), we draw, for a given sample $(x_i)_{1 \leq i \leq n}$, \mathcal{S}_p as a function of u (i.e. $\kappa = u$), observing if it admits a unique minimum. Take the example of a G-GPD sample $(x_i)_{1 \leq i \leq n}$ with $n = 10^3$, generated with $\mu = 0$, $\sigma = 1$ and $u = 0.5221$, and consider different scenarios varying the value of the Gaussian parameter p . In Figure 7, we present the different curves of \mathcal{S}_p depending on p , which all exhibit a unique minimum. Note that for $p = [0, 1]$ which corresponds to the parameter of the generated sample, the minimum of \mathcal{S}_p (see the green curve of the right plot) coincides with the exact threshold $u = 0.5221$ of the data (represented by a vertical red dashed line in both plots of Figure 7), as expected.

$$H(x; \theta) = \left[\frac{F(x; \mu, \sigma)}{1 + F(u; \mu, \sigma)} \right] \mathbb{1}_{\{x \leq u\}} + \left[1 - \frac{\left(1 + \left(\frac{u - \mu}{\sigma^2} - f(u; \mu, \sigma) \right) (x - u) \right)^{-\frac{f(u; \mu, \sigma)}{\frac{u - \mu}{\sigma^2} - f(u; \mu, \sigma)}}}{1 + F(u; \mu, \sigma)} \right] \mathbb{1}_{\{x \geq u\}}. \quad (\text{A.1})$$




Figure 7: \mathcal{S}_p as a function of the threshold u (x , μ and σ are fixed)

Using $\varphi_u : \mathcal{D}_p \rightarrow \mathcal{D}_u$ and $\varphi_p : \mathcal{D}_u \rightarrow \mathcal{D}_p$, the two steps of the first iteration of the algorithm can be given, for a fixed $\tilde{u}^{(0)}$, by the following relations:

$$\begin{cases} \tilde{p}^{(1)} = \varphi_p(\tilde{u}^{(0)}; x), \\ \tilde{u}^{(1)} = \varphi_u(\tilde{p}^{(1)}; x) = \varphi_u(\varphi_p(\tilde{u}^{(0)}; x); x). \end{cases}$$

More generally, for any $k \geq 1$, we can write

$$\tilde{u}^{(k)} = \phi(\tilde{u}^{(k-1)}; x), \quad (\text{A.2})$$

where the function ϕ is defined from \mathcal{D}_u to \mathcal{D}_u by: $\phi(u; x) = \varphi_u(\varphi_p(u; x); x)$.

Consequently, the algorithm can also be expressed as:

Algorithm 3 Iterative and unsupervised algorithm for the G-GPD parameters estimation
(version 2)

- 1: Initialization of $\tilde{u}^{(0)}$, $\epsilon > 0$, and k_{max} .
 - 2: Determination of the empirical cdf H_n according to x .
 - 3: Iterative process:

$k \leftarrow 1$

$$\tilde{u}^{(k)} = \phi(\tilde{u}^{(k-1)}; x)$$
 $k \leftarrow k + 1$
 while $|\tilde{u}^{(k)} - \tilde{u}^{(k-1)}| \geq \epsilon$ and $k < k_{max}$.
 - 4: Return $\tilde{u}^{(k)}$.
-

A way to prove the existence of stationary points of Algorithm 3 is to demonstrate the existence of fixed-points of the function ϕ . To do so, we build on the fixed-point theorem. Several versions of this theorem exist in the literature *e.g.* the version of Banach (see [57]), or of Markov-Kakutani (see [55]), or of Schauder (see [38]), or of Brouwer (see [24]). In this work, we consider the latter one, as its hypotheses are, in our case, more straightforward to check. This theorem states that *every continuous function from a closed ball of a Euclidean space into itself has a fixed point*. It implies that the functional ϕ admits at least one fixed point if the following two conditions, (\mathcal{C}_1) and (\mathcal{C}_2) , are satisfied:

(\mathcal{C}_1) : \mathcal{D}_u is a closed ball of a Euclidean space.

(\mathcal{C}_2) : ϕ is continuous on \mathcal{D}_u .

The conditions (\mathcal{C}_1) is clearly satisfied: for a realization x , $\mathcal{D}_u = [0, \max(x)]$ is a closed ball of \mathbb{R} that is a Euclidian space.

Now, to verify (\mathcal{C}_2) , we prove that φ_u and φ_p are both continuous on their domains (since ϕ is the composite function: $\phi = \varphi_u \circ \varphi_p$) using the Heine-Cantor theorem (see *e.g.* [50]) and the Ramsay et al.'s one that we recall here.

THEOREM ([45])

Let \mathcal{X} and \mathcal{Y} be metric spaces with \mathcal{X} closed and bounded. Let

$$\begin{aligned} g : \mathcal{X} \times \mathcal{Y} &\rightarrow \mathbb{R} \\ (x, \alpha) &\rightarrow g(x, \alpha) \end{aligned}$$

be uniformly continuous in x and α , such that $x(\alpha) = \underset{x \in \mathcal{X}}{\operatorname{argmin}} g(x, \alpha)$ is well defined for all $\alpha \in \mathcal{Y}$.

Then the function $x(\alpha) : \mathcal{Y} \rightarrow \mathcal{X}$ is continuous.

The proof of the continuity of the two functions φ_p and φ_u being the same, let us consider for instance the function φ_p . Using Ramsay's theorem, we need to check that \mathcal{D}_p is a compact and that \mathcal{S}_u is uniformly continuous on \mathcal{D}_p , to conclude to the continuity of φ_p . The first condition, \mathcal{D}_p is a compact of \mathbb{R}^2 , is satisfied when noticing that we are working with a Gaussian density, with finite mean and variance, hence which is bounded.

Now, as \mathcal{D}_p is a compact, it is sufficient to show that \mathcal{S}_u is continuous on \mathcal{D}_p to deduce, by the Heine-Cantor theorem, its uniform continuity. Since $\mathcal{S}_u(p; x) = \sum_{i=1}^n (H(x_i; \theta | u) - H_n(x_i))^2$, we

just need to study the continuity of H w.r.t. p to deduce the continuity of \mathcal{S}_u w.r.t. p . We recall here that, by construction, H is continuous w.r.t. x and not to its parameters. Hence, its continuity according to p remains to be proved. Since H is composed of two functions (see (2.3) for $u_1 = u_2$), the Gaussian cdf and the GPD, we will study the continuity of each one w.r.t. p . The continuity of the Gaussian cdf F as a function of p is immediate since it means to look at the continuity of its likelihood w.r.t. $p = [\mu, \sigma]$. Now, for the GPD G , its parameters ξ and β are expressed as functions of p : $\beta(p) = \frac{1}{f(u; p)}$, and $\xi(p) = \frac{u - \mu}{\sigma^2} \beta(p) - 1$, and are both continuous in p . Hence G is continuous in p as the composition of continuous functions w.r.t. p .

Finally, we can deduce the continuity of the function \mathcal{S}_u on \mathcal{D}_p as a composition, sum and products, of continuous functions on \mathcal{D}_p , from which we conclude to the continuity of φ_p on \mathcal{D}_u .

Conclusion: We can conclude that the functional ϕ is continuous on \mathcal{D}_u as a composition of two continuous functions: φ_p and φ_u . Hence the existence of at least one fixed-point according to the Brouwer fixed-point theorem. Consequently, the algorithm admits at least one stationary point. Since the method does not follow a path on an error surface, it is free from local minima traps as are the classical gradient search based methods. In the next section, we perform simulations to check if the algorithm converges to a unique stationary point regardless to the initialization.

B Numerical study of the algorithm convergence

To study numerically the convergence of the algorithm to a unique attractive stationary point, we consider the recurrent sequence $\{\tilde{u}^{(k+1)} = \phi(\tilde{u}^{(k)})\}_{k \in \mathbb{N}^*}$, obtained when applying Algorithm 2 on a generated G-GPD distributed data with a fixed parameter θ . Different initial values of this sequence are considered and for each one we represent graphically the associated recurrent sequence. To ensure the algorithm to be on the right track, all initial values are selected in the interval to which u belongs, namely $I = [q_{25\%}, q_{50\%}]$ (see [12]). For illustration, we report here two examples among all those performed to test the convergence.




Figure 8: Study of the convergence of the recurrent sequence $\{\tilde{u}^{(k+1)} = \phi(\tilde{u}^{(k)})\}_{k \in \mathbb{N}^*}$ regarding the initial value $\tilde{u}^{(0)}$. Example 1 for $\theta = [0, 1, 0.4354]$ and $u^* = 0.4354q_{39.42\%}$.

For $\theta = [0, 1, 0.4354]$ with $\xi = 0.2$ and $\beta = 2.7558$, we present in Figure 8 the recurrent sequence $\{\tilde{u}^{(k+1)} = \phi(\tilde{u}^{(k)})\}_{k \in \mathbb{N}^*}$, where the initial value $\tilde{u}^{(0)} \in \{q_{35\%}, q_{37.5\%}, q_{40\%}, q_{42.5\%}, q_{45\%}, q_{47.5\%}\}$. As shown in this figure, regardless the choice of $\tilde{u}^{(0)}$ in I , the algorithm converges to the fixed value of $u = 0.4354$ (represented by a continuous horizontal line), denoted by u^* . We observe that:

1. If $\tilde{u}^{(0)} < u^*$, the associated recurrent sequence is non decreasing, as for instance for the gray circles curve with $\tilde{u}^{(0)} = q_{35\%}$ and the red triangles (upwards oriented) one with $\tilde{u}^{(0)} = q_{37.5\%}$;
2. If $\tilde{u}^{(0)} > u^*$, the associated recurrent sequence is non increasing, *e.g.* the blue diamonds curve for $\tilde{u}^{(0)} = q_{45\%}$ and the pink triangles (downwards oriented) curve for $\tilde{u}^{(0)} = q_{47.5\%}$.

Consequently, based on Figure 8, regardless the choice of $\tilde{u}^{(0)} \in I$, the recurrent sequence $\{\tilde{u}^{(k+1)} = \phi(\tilde{u}^{(k)})\}_{k \in \mathbb{N}^*}$ is monotone on \mathcal{D}_u and converges to a unique attractive stationary point that corresponds to u^* .

Another example now for $\theta = [3, 2, 4.0443]$ with $\xi = 0.5$ and $\beta = 5.7454$ is illustrated in Figure 9 and leads to the same observations.




Figure 9: Study of the convergence of the recurrent sequence $\{\tilde{u}^{(k+1)} = \phi(\tilde{u}^{(k)})\}_{k \in \mathbb{N}^*}$ regarding the initial value $\tilde{u}^{(0)}$. Example 2 for $\theta = [3, 2, 4.0443]$ and $u^* = 4.0443 = q_{41.18\%}$

An additional remark concerns the number of iterations. We could observe in the simulation study that the closest to u^* is $\tilde{u}^{(0)}$, the fastest is the convergence, as expected. It appears clearly on the two reported examples (see the green '+' marks curve in both figures).

To conclude, let us comment that extending this convergence study to three-components is straightforward and follows the same logic as for two components. The estimation of θ is also broken down into the estimation of p and ξ alternately. The associated algorithm can be, as well as for the two-components model, represented by a functional of ξ . Hence, we can prove as previously that this functional is continuous on a closed ball of a Euclidean space, according to the Brouwer theorem, to infer the existence of stationary points of Algorithm 1. The difference between the two- and three-components algorithms concerns only the data scale, the stop condition, and the condition on ξ to be positive (Fréchet distribution), so does not interfere in the convergence of the three-components algorithm. These three conditions have only been introduced to enhance the parameters estimation.

PARIS

ESSEC Business School

3, avenue Bernard-Hirsch
CS 50105 Cergy
95021 Cergy-Pontoise Cedex
France
Tél. +33 (0)1 34 43 30 00
www.essec.edu

ESSEC Executive Education

CNIT BP 230
92053 Paris-La Défense
France
Tél. +33 (0)1 46 92 49 00
www.executive-education.essec.fr

ESSEC Asia-Pacific

5 Nepal Park
Singapore 139408
Tél. +65 6884 9780
www.essec.edu/asia

SINGAPOUR

Contact :

Centre de Recherche
+33 (0)1 34 43 30 91
research.center@essec.fr

ISSN 1291-9616



affilié à la
 CCI PARIS ILE-DE-FRANCE

