



HAL
open science

Flood frequency analysis: The Bayesian choice

Eric Gaume

► **To cite this version:**

Eric Gaume. Flood frequency analysis: The Bayesian choice. Wiley Interdisciplinary Reviews: Water, 2018, 11p. 10.1002/wat2.1290 . hal-01792478v1

HAL Id: hal-01792478

<https://hal.science/hal-01792478v1>

Submitted on 18 May 2018 (v1), last revised 22 Nov 2018 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Flood Frequency Analysis: the Bayesian choice

Eric GAUME*

Article Type: Primer

Abstract

After an introduction to the traditional flood frequency analysis methods, this article discusses their limits and the risks associated to their thoughtless use: overconfidence in the estimated values of flood quantiles or return periods and systematic underestimation of risks. The article then presents and illustrates the added value of modern Bayesian flood frequency inference procedures that are statistically consistent, numerically accurate and computationally now affordable. The implementation of such methods shows that estimated flood frequencies, based on observed samples of limited size, are generally affected by large uncertainties. This acknowledgement should be an incentive for increasing the size of the analysed samples through a more systematic use of historic information as well as regional approaches in flood frequency analyses. It also clearly points that margin of errors should be considered when using inference results for design or risk assessment purposes. Several pieces of software are now available to conduct Bayesian flood frequency analyses relatively straightforwardly. There is no remaining obstacle to the implementation of these modern approaches in operational hydrological studies.

*IFSTTAR, GERS, F-44344 Bouguenais, France

INTRODUCTION

Frequency analysis is a common engineering practice in many domains where risk assessment or design is required and singularly in hydrology. It has also been an active field of research and publication since the founding works of Fréchet (?, ?) and Fisher and Tippett (?, ?) on the theory of extreme values in statistics. Numerous methods have been proposed to extract samples of a variable of interest from observed data sets (sampling) and to fit statistical distributions to these samples (i.e. to infer the statistical properties of the underlying unknown natural random process). It is important to mention here that this paper is focussed on data-driven (i.e. inference-based) flood frequency analyses. Other methods, specifically developed to explore the tail of statistical distributions (?, ?) (i.e. provide reasonable values for very high return periods) will not be presented: (a) Probable Maximum Flood approach or (b) Derived flood frequency methods based on the coupling between a rainfall stochastic model and a rainfall-runoff model. The traditional inference procedures for flood frequency analysis and their background will be shortly presented and their respective merits discussed in the second section of this article, but it is not its main focus.

In practice, inference methods are often implemented like recipes, with the objective to provide an estimated quantile or return period value, corresponding to the statistical distribution best suited to the sample - the definition of the terms quantile and return period will be recalled hereafter. In fact, if limited to this optimisation objective, inference has little to do with statistics. In a consistent statistical perspective, the estimates, based on samples resulting from a random process, should also be considered as random variables. In other words, these estimates are variable or affected by uncertainties, related to sampling variability, that can and should also be evaluated and provided as part of the result of the inference procedure. Hiding estimation uncertainties, on purpose or by ignorance, gives a false impression of accuracy and rigorousness and proves Mark Twain right who stated ironically: *'there are three types of lies: lies, damned lies and statistics'*. Worse, it leads to biased decisions, to underrate the risks in fact, as will be illustrated herein.

This article therefore presents and clearly aims at promoting Bayesian frequency analysis methods, that are theoretically consistent, have proven to be numerically accurate and are

now computationally affordable. The title ‘Bayesian choice’ is a wink to the famous textbook written by Christian P. Robert (?). Engineers may not feel at ease with inference results affected by large ‘posterior’ uncertainties (i.e. estimated values with large variances). They may even be reluctant to reveal these uncertainties. But it is the simple state of facts that estimated values based on limited samples are uncertain. Its acknowledgement should be an incentive to increase the size of the analysed data samples, taking advantage of historic information or of statistical homogeneities that may exist at a regional scale.

Many pieces of software are now available to conduct Bayesian flood frequency analyses. The illustrations of this paper are based on the freely available R software library nsRFA and more precisely on the BayesianMCMC command. There is no remaining obstacle for the implementation of such approaches in operational hydrological studies. But the readers must also be aware that even if Bayesian approaches constitute undoubtedly a progress in flood frequency analyses, they have also limits and some strong hypotheses are needed to conduct inferences such as the choice of a statistical distribution type. Bayesian approaches should not be the new god to be worshipped. No model, no method compensate for a lack of data. This essential idea is recalled in the conclusion of the paper: accuracy can only be achieved in flood frequency analyses if they are based on good-quality but also rich datasets. Data is the key factor, the crux of every statistical analysis.

The article is divided in two parts. The first part exposes the standard flood frequency inference procedures. The second part is devoted to uncertainties related to sampling variability and their consequences and presents the Bayesian inference method. A short conclusion recalls the importance of data in statistical analyses.

PART 1: BASICS OF FLOOD FREQUENCY ANALYSES

Data critical analysis

In flood frequency analyses, discharge is generally preferred to river stages. Discharges from different sites can be compared which enables cross-checking and regional frequency analyses. Water levels are highly dependent on the local flow conditions that may even vary

significantly over time, leading to non-stationary and inconsistent data series.

Discharge is nevertheless seldom directly measured but estimated based on measured stages and calibrated stage-discharge relations. The largest discharge values in a series, that are of crucial importance for flood frequency analyses, often correspond to the extrapolated and hence less accurate part of the stage-discharge relation.

It is therefore highly recommended to proceed to a detailed critical analysis of the data sets before any further computation. A particular attention must be given to:

- Data completeness: are the existing gaps not corresponding to possible large floods, that may for instance have damaged the stream gauge ? This can be verified looking at rainfall or nearby stream records.
- Data consistency: is the stage-discharge relation reliable ? Has the same relation been used over the whole record period or, if not, are the modifications justified ? It is frequent that hydrometric services recalibrate the stage-discharge relation to account for new measured (gauged) discharges. This generally has little impact on the interpolation range of this relation but may sometimes lead to major and unjustified evolutions of its extrapolated part. Such procedures may lead to highly inconsistent discharge series (see Figure ?? for an illustration).

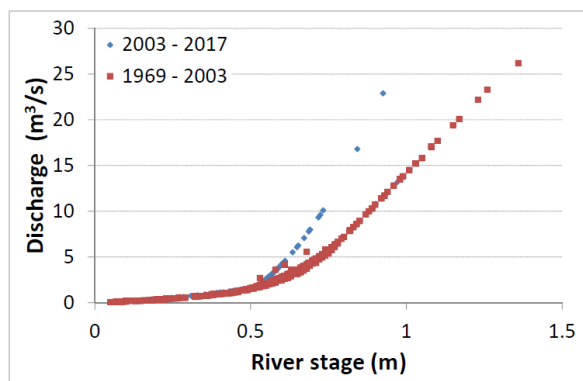


Figure 1: Plot of measured stages versus corresponding estimated discharges at the Montcel gauge on the Morge river (France) according to the available data set (left panel). View of the Morge river in the vicinity of the gauged cross-section (right panel). The stage-discharge relation has obviously been modified in 2003 with little changes in its lower part but tremendous evolutions of its upper, extrapolated, part.

Data sampling strategies

Flood frequency analyses aim at deciphering the natural random process driving the occurrence and magnitude of flood events. The objects of interest are flood events: i.e. periods of temporary increase of the river discharge above its average value. An event descriptive variable has to be selected to rate the magnitude of a flood. This variable may vary depending on the objective of each specific study. The peak discharge is often chosen. But other variables may be considered, such as the flood volume over a given discharge threshold for the design of dam volumes for instance.

Floods and their corresponding descriptive variable, once selected, have then to be sampled from the measured sequential series. Two main sampling strategies exist, illustrated in Figure ?? : (a) block maxima or (b) peaks over threshold (POT) sampling. The samples resulting from the two methods may have the same size, the same largest values, but are clearly different as illustrated in Figure ??.




Figure 2: Illustration of the two classic flood peak discharge sampling strategies applied to a ten year series of discharges measured in Paris on the Seine river: annual block maxima (left) and peaks over threshold (right). Red dots (sampled values), dotted vertical lines (separation between hydrological years), red horizontal line (selected threshold).

The block maxima sampling method is straightforward. The largest values observed during each available complete hydrological year (i.e. starting in October) have been for instance selected in the example shown in figure ??a. The POT method is probably at first sight the most natural one: every event exceeding a threshold value is a flood event and should be considered. It is the method recommended in some textbooks (?, ?). But it has also some drawbacks. First, the selection of floods remains partly arbitrary. Two successive periods of exceedance may be considered as belonging to the same event if they are not separated by a sufficiently long period of time. A minimum delay between two successive peaks is generally defined in POT sampling. Second, the threshold and hence the sampling strategy is site dependent. This adds a source of variability in regional flood frequency analyses or if quantiles computed at different sites are compared. For the simplicity and reproducibility of its implementation, the author prefers the annual maxima sampling method.

UNIVARIATE OR MULTIVARIATE STATISTICS

Some readers will be surprised that the manuscript is limited to univariate statistical methods when multivariate approaches have been flourishing in the recent years in the scientific literature, especially with the introduction of copulas in hydrology. But the notions of frequency or return period is only simply defined in the univariate case.

Multivariate statistical models may be developed to describe the joint probability of two or more flood descriptive variables (volume, duration, peak discharge...). Multiple definitions of an event are then possible, based on combinations of these descriptive variables or on their consequences (i.e. possible resulting water levels). Additional steps, generally based on Monte Carlo random simulations are then needed to provide estimates of return periods or exceedance probabilities of the defined events.

Multivariate approaches pertain to stochastic modelling rather than directly to frequency analysis.

Selection of a distribution type

Next step, a family of mathematical function has to be selected as candidate for the cumulative distribution of the sampled random variable X : $F_\theta(x) = P(X \leq x)$. θ stands for the set of parameters to be calibrated against the sample. F_θ can be any non-decreasing and right-continuous function such as $\lim_{x \rightarrow -\infty} = 0$ and $\lim_{x \rightarrow \infty} = 1$. Some families, with well-established mathematical formulations are popular in flood frequency analyses, such as the 3-parameter gamma (also called Pearson III), log-normal and log-Pearson III (the logarithm of X has a normal or Pearson III distribution), or one the three extreme value (EV) distributions (Gumbel or EV I, Fréchet or EV II and Weibull or EV III) or even the synthetic form of these three distributions, established by Gnedenko in 1943 (Gnedenko, 1943) and called General Extreme Value (GEV) distribution:

$$F_\theta(x) = \exp \left[- \left(1 - \frac{k(x-a)}{b} \right)^{1/k} \right]_{b>0} \quad (1)$$

In equation ??, the parameter set θ stands for $\{a, b, k\}$, with a a position parameter, b a scale parameter and k the shape parameter ($k = 0$ corresponds to the Gumbel, $k > 0$ to a Weibull and $k < 0$ to a Fréchet distribution).

It is time for some definitions. The so-called return period $T(x)$ of the value x is equal to : $T(x) = 1/(1 - F_\theta(x))$. T is expressed directly in years if annual maxima have been sampled. A quantile x_T is the value of the random variable X corresponding to return period T : $x_T = F_\theta^{-1}(T)$. the probability density function $f_\theta(x)$ is the partial derivative of F with respect to x .

$$f_\theta(x) = \frac{dF_\theta}{dx}(x) \quad (2)$$

DEFINITION OF THE RETURN PERIOD

Even if not appropriate, the term 'return period' T has become popular among experts as well as the general public. It is a source of misconception: the studied stochastic processes are by no way periodic! Moreover, several definitions of T coexist, that are nevertheless basically equivalent. T can be defined as the inverse of the annual probability $1 - F$ of exceeding a given quantile value x_T . Observed series of maximum discharges are generally supposed to be sequences of independent and identically distributed variables. In that case, the sequence of exceedances (success) and non-exceedances of x_T is a Bernoulli process, with success probability $1 - F$. The number of trials (number of years) N needed to get one success (exceedance) has a geometric distribution with expectancy $E[N] = 1/(1 - F) = T$. T is then likewise the expected waiting time before the next occurrence of a value larger than x_T . It does not depend on the past, since the successive events are independent. T is therefore also the expected duration between to successive exceedances. It as been shown that these equivalences resist to some level of dependence (?, ?). Note that the median waiting time is equal to $-\log(2)/\log(F)$ (i.e. 69 years for $T = 100$ years). And, the number of exceedances over a period of time M has a binomial distribution $B(M, 1 - F)$. The probability of observing at least one event equal or larger than the 100-year event over a 100-year period is equal to 63.4%...

Various tests have been proposed to evaluate the adequacy between a sample and a distribution type and to facilitate the choice of the distributions to be calibrated (?, ?). The GEV distribution plays nevertheless a particular role in flood frequency analysis. The GEV

distribution for the extreme value theory is an equivalent of the normal distribution for the central limit theorem. It is the limit distribution of the maximum value of n identically distributed random variables when n tends to infinity. This explains why, extreme value distributions are often chosen as candidate distributions in flood frequency analyses. But, it must be clear that the extreme value theory is an asymptotic theory. The distribution of observed samples of limited size can differ significantly from the GEV distribution. This is clearly acknowledged by Coles (1998, 2001) in the introduction of his famous textbook: *‘In Lieu of an empirical or physical basis, asymptotic arguments is used to generate the extreme value models. It is easy to be cynical about this strategy, arguing that extrapolation of models to unseen levels requires a leap of faith, even if the models have an underlying asymptotic rationale. There is no simple defense against this criticism, except to say that applications demand extrapolation, and that it is better to use techniques that have a rationale of some sort’.*

Hereafter, we will be working with GEV distributions, not merely because it is related to extreme value theory but mainly because it is a 3-parameter and hence flexible distribution type, leaving more space for data driven results.

Inference procedures

The last step of the analysis consists in estimating (inferring) the parameter values of the selected distribution or distributions best suited to the sample. Before proceeding to the inference, it is recommended to verify, through statistical tests, that the sample can reasonably be considered composed of independent and identically distributed values. Non-stationary statistical models may be proposed if clear temporal trends in the sample mean or variance are detected for instance. The presentation hereafter is limited to the standard stationary case study.

Three inference methods are proposed in the literature:

1. **Method of moments.** The moments (expectancy μ , variance σ^2 , skewness s and kurtosis k coefficients) and the parameters θ are related for a given distribution type: $\theta = g(\mu, \sigma^2, s, k)$. The empirical moments of a sample can be considered as estimators

of the moments of the underlying distribution $(\hat{\mu}, \hat{\sigma}^2, \hat{s}, \hat{k})$. Hence the same relations can be used to provide estimators of the parameters $\hat{\theta} = g(\hat{\mu}, \hat{\sigma}^2, \hat{s}, \hat{k})$.

2. **L-moments method.** The L-moments method is similar to the method of moments, but is based on other types of moments, the L-moments, introduced by S. Hosking in 1990 (?). L-moments are linear combinations of order statistics. L-moments have the advantage to be less sensitive than moments to sampling variability and have therefore become popular for flood frequency analyses, especially suited, at least theoretically, for 3-parameter distributions and small sample size.
3. **maximum likelihood.** The distribution type being selected and the sample being supposed to be composed of independent and identically distributed variables, the likelihood of the sample can be computed as the product of the densities of each item of the sample (standard definition, left side of equation ??) or the product of the probability of each item (definition preferred by the author, right side of equation ??). x_i^h and x_i^l represent respectively the expert-based high and low bounds of the confidence intervals for the estimated discharges. The values of the parameter set θ maximizing the likelihood $\mathcal{L}(X|\theta)$ of the sample are selected as estimators.

$$\mathcal{L}(X|\theta) = \prod_{i=1}^s f_{\theta}(x_i) \quad or \quad \mathcal{L}(X|\theta) = \prod_{i=1}^s [F_{\theta}(x_i^h) - F_{\theta}(x_i^l)] \quad (3)$$

Once the parameters $\hat{\theta}$ estimated, any quantile value can be evaluated: $\hat{x}_T = F_{\hat{\theta}}^{-1}(T)$.

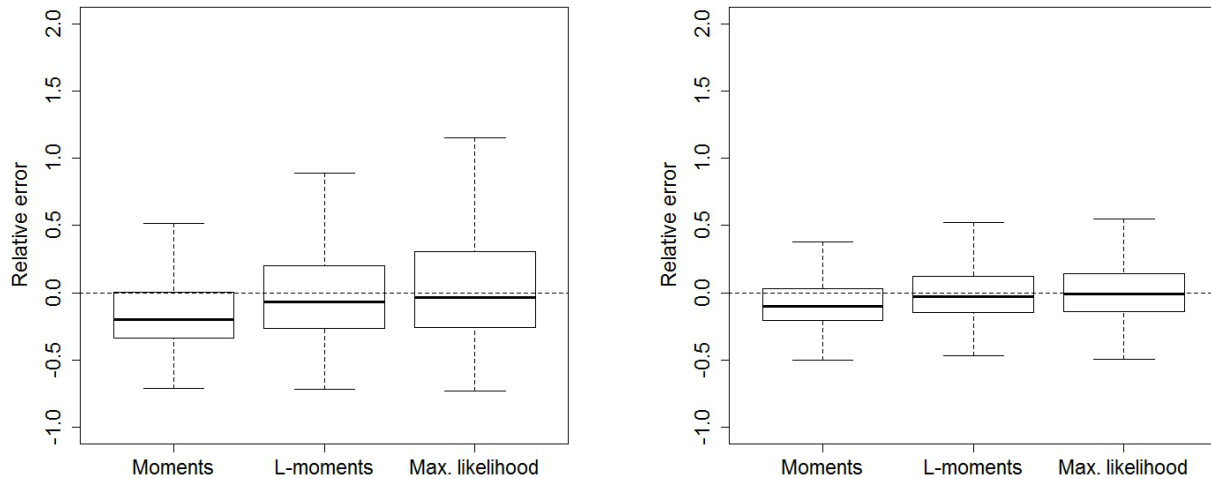


Figure 3: Boxplots of the 100-year quantile estimation relative errors for the three standard inference methods. 1000 series of N values drawn from a parent $GEV(10,10,-0.2)$ distribution. $N = 30$ (left) and $N = 100$ (right).

Figure ?? illustrates the performance of these three inference methods when applied to 1000 samples, randomly drawn from a $GEV(a,b,k)$ distribution: relative estimation error for the 100-year quantile $(\hat{x}_{100} - x_{100})/x_{100}$. No general conclusion can be drawn from one single example, but the maximum likelihood is known to be theoretically unbiased (i.e. with expected estimation error equal to zero), but biases may nevertheless be introduced by the numerical optimization method needed. It is also asymptotically optimal (i.e. with minimum error variance). Without surprise, the estimator based on maximum likelihood appears almost unbiased (median as well as mean error value close to zero), whereas estimators based on moments are biased and have a tendency to under-estimate the 100-year quantile for this specific distribution. The maximum likelihood estimator has the largest error variance for the small sample size (30 values in Figure ?? left) but this variance decreases rapidly with sample size. It is asymptotically the estimator with the lowest variance. More surprising, the variance of the L-moments estimator is larger than the variance of the estimator based on moments. L-moments are less sensitive to sampling variability, but parameters and quantiles are related to the moments by non-linear functions. The advantage of the lower variance of

the L-moments may be lost due to this non-linear transformation for this specific distribution.

Figure ?? also shows that quantile relative estimation errors may be large when the size of the available data sets is limited: often 50% or more. These errors, mainly related to sampling variability, will be the main focus of the next part of this paper.

PART 2: BAYESIAN INFERENCE

Sampling variability and uncertainties

The statistical distributions ($x_T = F_\theta^{-1}(T)$) corresponding to one of the samples used to build Figure ?? are drawn on the left side of Figure ??: the parent GEV(10,10,-0.2) as well as the three adjusted distributions based on the moments, L-moments and maximum likelihood. The sampled values with their corresponding empirical return periods (plotting positions) are added on the figure. The empirical return period is computed here in the following way. The sample of size M is ranked in descending order. The empirical return period T_i of the rank i value is then set equal to $T_i = (M + 1)/i$.¹

¹Note, that the empirical return periods are only used to enable a graphical comparison between the sample and the adjusted distributions, but were not use in the inference process. Accurate and especially unbiased plotting positions are not really needed. Several other formulas have been proposed for plotting positions in the past (?, ?) when some statistical analyses were based on graphical interpolations of plotting positions, an approach now abandoned.

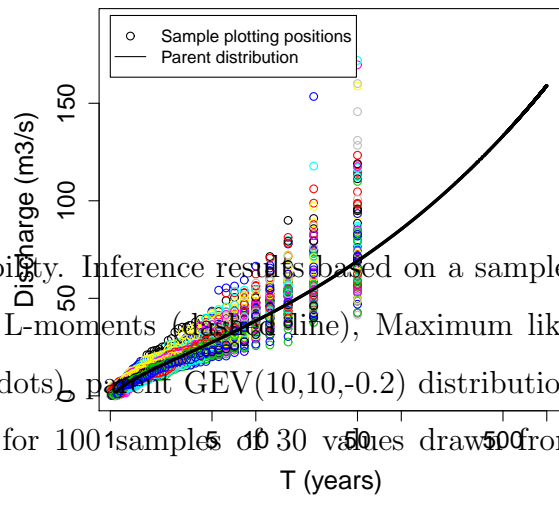
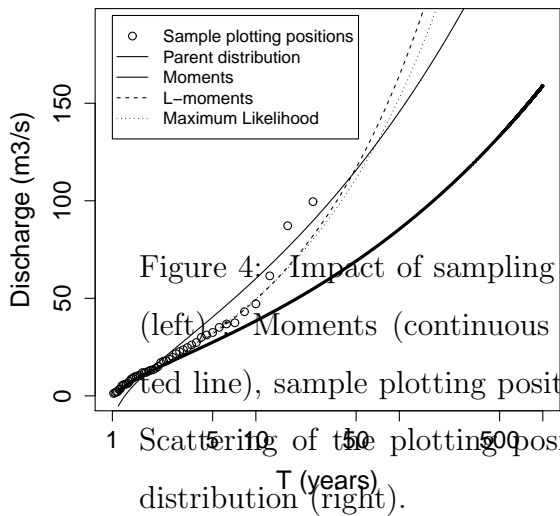


Figure 4: Impact of sampling variability. Inference results based on a sample of 30 values (left): Moments (continuous line), L-moments (dashed line), Maximum likelihood (dotted line), sample plotting positions (dots), parent GEV(10,10,-0.2) distribution (bold line).
 1. Scattering of plotting positions for 100 samples of 30 values drawn from the parent distribution (right).

It is clear from Figure ?? that the three calibrated distributions may be different, but are firstly controlled by the sample. This sample may drive the calibration far from the parent distribution (bold line on the figure). Of course, this example has been selected for the sake of explanation. But Figure ?? in its right panel, illustrating the scattering of plotting positions due to sampling variability, shows that this specific sample of 30 values is not unusual for such a GEV parent distribution. Clearly, the choice of an inference method is a problem of secondary importance in flood frequency analysis. The main problem is the important effect of sampling variability. This being acknowledge, three main questions will be discussed hereafter:

1. Is it possible to assess inference uncertainties due to sampling variability when only one sample, the measured sample, is available and how ?
2. What are the possible solutions to increase the size of the studied samples in order to reduce the effects of sampling variability ?

3. How can inference uncertainties be accounted for in a design or risk assessment study ?

Estimation of inference uncertainties: bootstrap and Bayesian inference

The sample of M items entails information about the variability of the underlying random process. To evaluate the effect of sampling variability on flood frequency analyses, a first idea consists in re-sampling randomly, with replacement, from the available data set, numerous samples of M values. A frequency analysis is conducted on each of these synthetic samples to assess the resulting variability and distributions for the parameters $\hat{\theta}$ and the corresponding quantiles \hat{x}_T . This empirical procedure is called bootstrap and is a standard method for assessing the properties of estimators.

An alternative approach can be proposed to evaluate the possible distribution of the parameters θ given the observations X . According to the Bayes theorem, the conditional probability density function $p(\theta|X)$ is related to the previously calculated likelihood $\mathcal{L}(X|\theta)$ (equation ??):

$$p(\theta|X) = \frac{\mathcal{L}(X|\theta)p(\theta)}{p(X)} \quad (4)$$

In equation ??, $p(\theta)$ is called prior distribution of θ . It summarizes any prior or alternative knowledge on θ . $p(X)$ is the probability of the data sample X which is unknown but fixed. When no prior information exists on θ (i.e. non-informative prior), $p(\theta)$ can be taken equal to 1, one common choice among other alternatives. This implies that $p(\theta|X)$ is proportional to $\mathcal{L}(X|\theta)$.

Numerical methods combining Monte Carlo simulations and Markov Chain random walk (MCMC) have been developed to explore the hyperspace of variables (θ in the present case) according to their joint probability. They are a class of algorithms for sampling from multivariate random distributions whose density is known ($?, ?, ?$): i.e. draw samples of parameter sets θ according to the density $p(\theta|X)$ in the present case. Typically, tends to hundreds thousands sets are sampled to be able to accurately describe the posterior joint distribution of the parameters and the corresponding distributions of quantiles. A proportionality constant

does not influence the result: i.e., the computation of $\mathcal{L}(X|\theta)$ or $\mathcal{L}(X|\theta)p(\theta)$ if a prior distribution is selected, is sufficient to sample vectors θ from a distribution with density function $p(\theta|X)$, without any knowledge of $p(X)$.

The final result of the Bayesian-MCMC procedure is a set of sampled vectors θ , typically some ten thousand vectors, with density $p(\theta|X)$. The corresponding discharge quantiles can then be computed and uncertainty bounds, credibility intervals using the Bayesian vocabulary, estimated based on this large set of quantile values.

In the Bayesian perspective, the result of the inference is not an optimum vector of parameters θ , but a density function $p(\theta|X)$ (conditional to the data set, i.e. 'posterior'). If the inference problem is well-posed, the spread of this posterior distribution should diminish as the information content of the sample X increases: i.e., this spread can be seen as a measure of the information content of the sample.

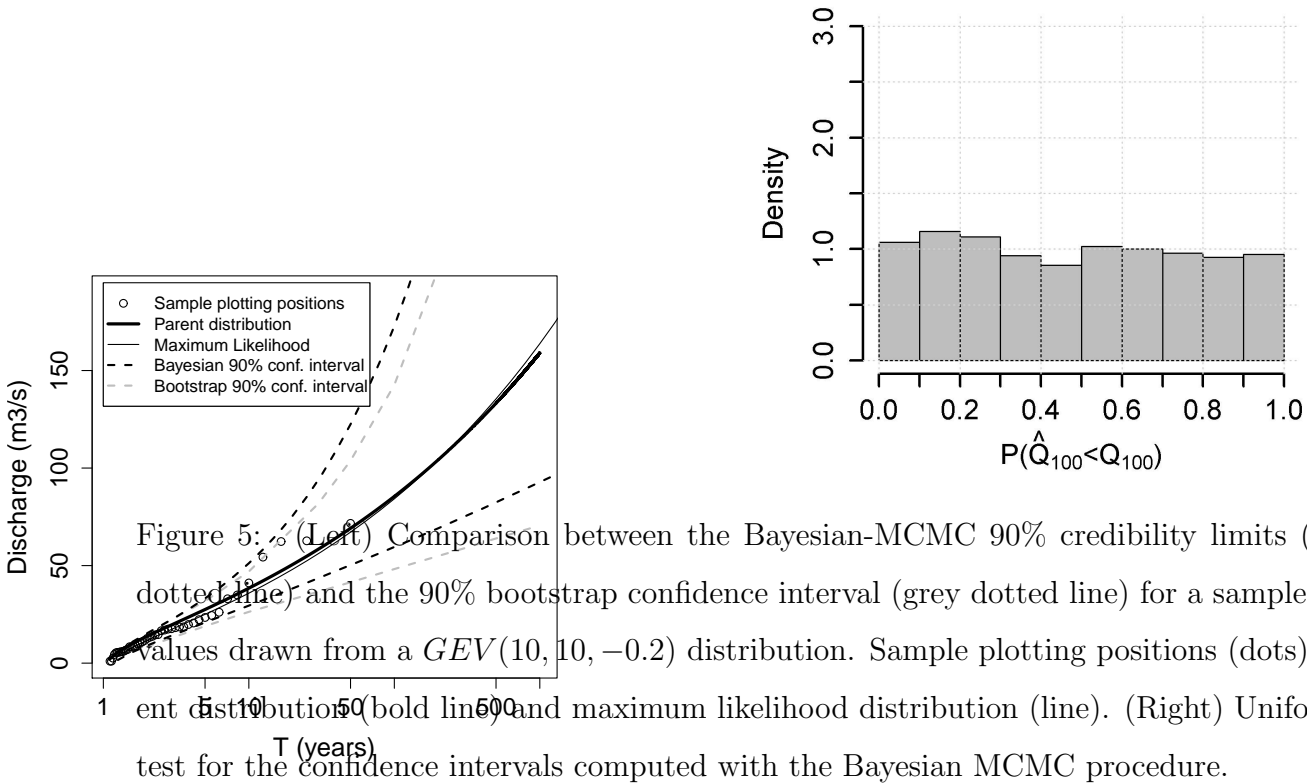


Figure 5: (Left) Comparison between the Bayesian-MCMC 90% credibility limits (black dotted line) and the 90% bootstrap confidence interval (grey dotted line) for a sample of 30 values drawn from a $GEV(10, 10, -0.2)$ distribution. Sample plotting positions (dots), parent distribution (bold line) and maximum likelihood distribution (line). (Right) Uniformity test for the confidence intervals computed with the Bayesian MCMC procedure.

Figure ?? compares, for one sample drawn from a $GEV(10, 10, -0.2)$ distribution, the

two 90% confidence intervals for the inferred quantiles computed with bootstrap (dotted grey line) and the Bayesian-MCMC approach (dotted black line). The two confidence intervals are comparable even if not totally super-imposed. The right side of Figure ?? shows the result of an accuracy test conducted on the posterior distributions of the \hat{x}_{100} quantile computed with the Bayesian-MCMC method (?, ?). It represents the distribution of the non-exceedance probability of the real quantile value x_{100} according to the posterior distribution of \hat{x}_{100} . If this posterior distribution is unbiased, the non-exceedance probability should be uniformly distributed when computed over a large amount of generated samples. It seems to be the case. This confirms that, even if the confidence intervals may appear large on figure ??, they are accurate, provided that the studied random variable is distributed according to the tested distribution. This is never sure in real case studies and the estimation uncertainties should be considered as underrated, even if large, since they have been computed conditionally to a distribution type. These large intervals reflect the limited information content of the available data set.

Reducing estimation uncertainties

Two main possibilities exist to enrich a data set, i.e. increase its size: 1) increase the observation time frame using existing information on historic floods and 2) take advantage of statistical homogeneities at a regional scale to built regional samples. Both approaches may also be combined (?, ?).

Valuation of historic information

Historic information has a particular format. Information, generally maximum water stages or flood extent maps, is only available for the largest flood events, flood events that have produced inundations and damages, which peak discharge value has exceeded a given perception threshold y_P . Historic information produces a censored data set Y . But censored information can straightforwardly be introduced in likelihood formulations (?, ?) (equation ??) and the previously presented inference procedures can be implemented without any further difficulties.

$$\begin{aligned}
\mathcal{L}(X, Y|\theta) = & \underbrace{\left[\prod_{i=1}^M [F_\theta(x_i^h) - F_\theta(x_i^l)] \right]}_{(a)} \cdot \underbrace{\left[\prod_{k=1}^{h_1} F_\theta(y_k^h) - F_\theta(y_k^l) \right]}_{(b)} \\
& \cdot \underbrace{\left[(1 - F_\theta(y_P))^{h_2} \right]}_{(c)} \cdot \underbrace{\left[F_\theta(y_P)^{h_3} \right]}_{(d)} \tag{5}
\end{aligned}$$

In equation ??, the term (a) is the same as in equation ?? and corresponds to the continuous series of measured data. The term (b) is the likelihood corresponding to the h_1 historic floods for which a discharge estimate is available. The term (c) represents the likelihood of the h_2 historic floods for which no other information is available than a given threshold y_P has been exceeded. The last term (d) is essential and should not be forgotten, it corresponds to all the h_3 years in the historic period where no record is available because the threshold y_P has most probably not been exceeded. The likelihood formulation may even account for successive historic periods with various perception threshold values if needed.

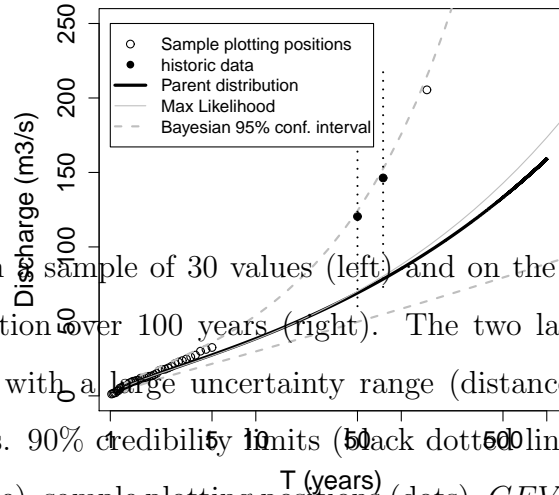
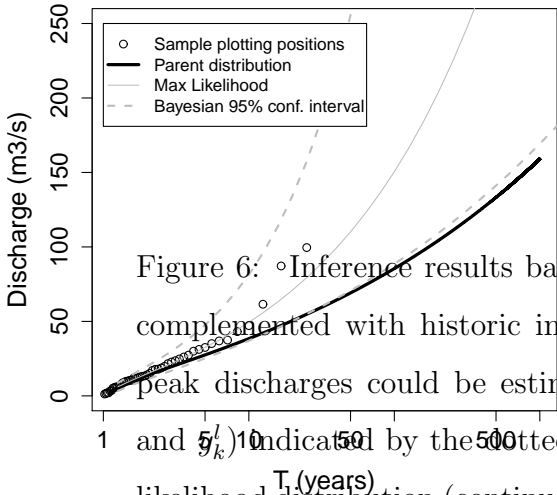


Figure 6: Inference results based on a sample of 30 values (left) and on the same sample complemented with historic information over 100 years (right). The two largest historic peak discharges could be estimated with a large uncertainty range (distance between y_k^h and y_k^l) indicated by the dotted lines. 90% credibility limits (black dotted line), maximum likelihood distribution (continuous line), sample plotting positions (dots), $GEV(10, 10, -0.2)$ parent distribution (bold line), historic records (vertical dotted lines). Note that the plotting positions have been recomputed in the right panel to account for the historic records.

Figure ?? gives an illustration of the possible added value of historic records. The two largest discharges over a historic period of 100 years could be estimated with large uncertainties, complementing a series of 30 years of continuous records. The threshold y_P has been set equal to the smallest y_k^h value: i.e. it is essential to be certain that the threshold has not been exceeded during the remaining historic years; its value should not be set too low. The same critical attention must be given to historic data as to measured records even if discharge estimates are less accurate. Misuse of historic information can result in significant flood frequency analysis biases. First, even if affected by large uncertainties, discharge estimates in the historic periods must be consistent with recent records. Second, it is essential that the historic inventory of events exceeding the perception threshold is exhaustive. If a doubt exists, the threshold has to be increased to limit the risk of an unreported flood

event exceeding it. Finally, the length of the historic period ($h_1 + h_2 + h_3$) has to be wisely selected based on all available information. It generally begins significantly before the first reported large flood event. If this event has drawn attention, it is generally because it appeared unusual if compared to previously observed events. Provided that this conditions are met, historic information generally help reduce significantly inference uncertainties even if historic discharges can only be roughly estimated. According to recent studies, the length of the historic recording period is much more important for frequency analyses than the accurate retrieval of historic peak discharges (?, ?).

Regional flood frequency analyses

Another possible option to increase the size of the data sets consists in merging observations available at different sites to build statistically homogeneous regional data sets. Regional frequency analyses is not a new idea. It is, for instance at the core of the work conducted by Hershfield and Kohler in 1960 (?, ?) which aimed at testing the Gumbel distribution against the daily rainfall series observed over the whole area of the United States. But it has received a renewed attention since the publication of book of Hosking and Wallis in 1997 on the topic (?, ?). This manual is the origin of the nsRFA library. In the simplest version of regional flood frequency analyses, it is hypothesized that in an homogeneous region, the local distributions are the same provided that the local values are scaled by a local constant factor μ^i called 'index flood' : $x_T^i = \mu^i \chi_T$ where i is the index of site, x_T^i the local and χ_T the regional quantiles.

In the case of a GEV distribution, it implies in particular that the distributions share the same shape parameter. Of course, the plausibility of this hypothesis should be statistically tested before conducting any inference on the regional composite data set. The 'index flood' is proportional to the expectancy of each local distribution. It is generally approximated by the average or the median value of each local sample.

In most cases, regional frequency analyses are implemented to merge observed data sets from different gauged sites. But the method has been recently successfully extended to account for discharges of extreme floods estimated at ungauged sites in a given region (?, ?). This opens new perspectives for the exploration of the possible shape of discharge distribu-

tions for return periods ranging from 100 to 1000 years.

Uncertainties and decision making

It is a common practice to take a margin of error for design purposes when computations are affected by possible uncertainties. Is it then reasonable to select the T -year quantile corresponding to the maximum likelihood for a T -year failure design and risk assessment in hydrology, since its estimation is affected by large uncertainties? The general answer is negative. The objectives set to the design have to be considered for the selection of the proper design value z_T based on the statistical inference result. If the risk of undersizing has to be limited, for instance to 10%, the upper bound of the T -year quantile 80% credibility interval has to be selected as design value. An other common objective is to select a design value that has an expected probability of exceedance equal to $1/T$ according to the outcome of the inference procedure including uncertainties. It may be surprising, but this design value is not the T -year maximum likelihood quantile \hat{x}_T as already shown by Stedinger in 1983 (Stedinger, 1983) and illustrated in figure 10.1.a.

In figure 10.1.a, as in the previous examples, 2000 samples of 30 values were drawn from a GEV distribution. The 2000 corresponding 100-year maximum likelihood estimated quantiles \hat{x}_{100} have been computed. The distribution of the exceedance probabilities $1 - F(\hat{x}_{100})$ of these estimated quantiles, according to the parent distribution, is plotted in the histogram. The average probability of exceedance of the quantile \hat{x}_{100} appears to close to 2% instead of 1%. The expected probability of exceedance (i.e. probability of failure of the designed structure) of the \hat{x}_T maximum likelihood estimated quantile is larger than $1/T$. Its selection as design value will lead, on average, to too frequent failures of the designed structures.

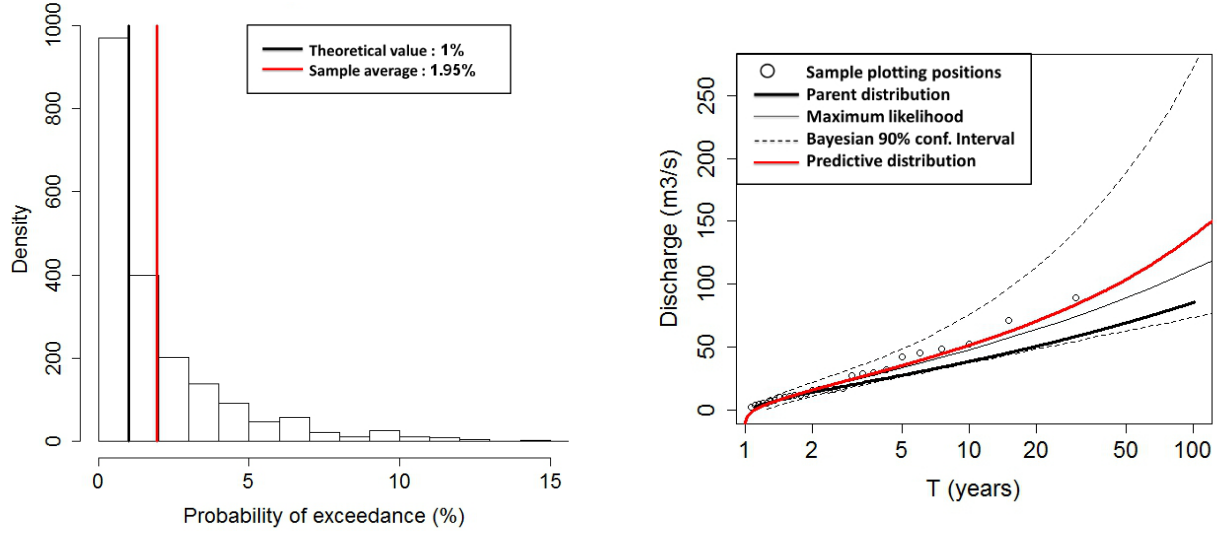


Figure 7: (left) Distribution of exceedance probabilities ($1 - F(\hat{x}_{100})$) of the 100-year estimated quantiles \hat{x}_{100} for 2000 samples of 30 values drawn from a $GEV(10, 10, -0.2)$ distribution: red vertical line (theoretical value: 1%), black vertical line (sample average). (right) Illustration of the 'predictive' distribution in red along with parent distribution (bold line), maximum likelihood distribution (continuous line), 90% confidence interval (dotted lines), sample plotting positions (dots).

The theoretical distribution of the design values z_T with the expected annual probability of failure $1 - E[F(z_T)] = 1/T$ (red line on the right panel of figure ??), is sometimes called 'predictive' distribution (?, ?). The function $E[F(z)]$ can be computed based on the output of the Bayesian inference procedure and especially on the posterior distribution $P(\theta|X)$ for the parameters θ : $E[F(z)] = \int_{\theta} F_{\theta}(z)P(\theta|X)d\theta$. Practically, this expectancy is estimated by the average value resulting from a sample Θ of M parameter sets θ randomly drawn in the posterior distribution $P(\theta|X)$: $E[F(z)] \approx \frac{1}{M} \sum_{\theta} F_{\theta}(z)$. Note that the 'predictive' distribution is numerically estimated. Its type and the mathematical form of its density function are unknown and probably differ from the GEV distribution type. This being said, it may appear complex to compute the 'predictive' distribution in operational applications. Moreover, it should be kept in mind that uncertainties related to the fact that the real distribution of the observed variables is unknown are not accounted for. The selection of

a quantile value for design purposes thus often consists in selecting the upper bound of a computed posterior confidence interval for the selected quantile - typically the 70% interval.

IMPORTANT CONCLUDING REMARK

The introduction of Bayesian inference methods in flood frequency analyses is undoubtedly a major progress, but their users should also be aware of their limits. The numerical derivations necessitate the selection of a theoretical mathematical function - a statistical distribution - to be fitted to the data sets. The inference results are constrained by this a priori and adhoc choice. But the natural processes seldom fit perfectly to theoretical mathematical models. As recalled by the mathematician Steward, cited by Klemes (1988) (? , ?):*'We may write down equations, and nature may - at some level - obey them, but nature is not obliged to restrict herself to those solutions that our overgrown monkey intellects can write down explicitly. And so mathematicians must pay attention to what really happens, rather than assume that nature conspires to make human calculations easy'*.

Besides, recent research results suggest that natural distributions of flood peak discharges may depart significantly from theoretical models and especially the asymptotic GEV model (? , ? , ?). Trends for medium and large return periods may not be captured in short duration observed series. Even if already large, the uncertainties affecting the inference results may therefore be significantly under-estimated.

No model and no method can replace data in flood frequency analyses. As mentioned in introduction, data is the key factor, the crux of every statistical analysis. Engineers and researcher should therefore focus their attention in the future on the careful documentation of extreme historic or regional events, even events that have occurred at ungauged sites. Increasing information on large and extreme events is the only possible way to gain knowledge about the tail of statistical distributions and reduce the uncertainties related to flood frequency analyses.

ACKNOWLEDGEMENTS

My thanks to Jose-Luis Salinas, Elena Volpi and Alberto Viglione whose thorough reading and detailed feedback helped to significantly improve the first version of this manuscript.

FURTHER READING

The nsRFA library of the free software R and its documentation provide helpful functions and explanations for readers who want to rapidly be able to implement Bayesian flood frequency inference methods.