

Discrete Extremes

BY A. S. HITZ

*Department of Statistics, University of Oxford,
 24-29 St Giles, Oxford OX1 3LB, UK
 adrien.hitz@stats.ox.ac.uk*

5

R. A. DAVIS

*Department of Statistics, Columbia University,
 1255 Amsterdam Avenue, New York, NY
 rdavis@stat.columbia.edu*

AND G. SAMORODNITSKY

*School of Operations Research and Information Engineering, Cornell University,
 220 Rhodes Hall, Ithaca, NY.
 gs18@cornell.edu*

10

SUMMARY

We propose two peaks-over-threshold methods for discrete random variables, and show that they can provide accurate tail probability estimates in simulated and real data.

15

Some key words: Extreme value theory; Tail approximation; Count data; Discrete distribution; Generalized Pareto distribution; Zipf distribution; Peaks over threshold.

1. INTRODUCTION

Extreme quantile estimation is an important but difficult problem in statistics, especially when the quantile is beyond the range of the data. In the univariate case, an approach that often works well in practice is to model observations above a large threshold with a parametric family of distributions which can be motivated as follows. Let X be a random variable taking values in $[0, x_F)$ for $x_F \in (0, \infty]$, and suppose that there exists a strictly positive sequence a_u such that

20

$$a_u^{-1}(X - u) \mid X \geq u \rightarrow Z, \quad (1)$$

25

in distribution as $u \rightarrow x_F$, for some Z following a non-degenerate probability distribution on $[0, \infty)$. Then, Z follows a generalized Pareto distribution, defined by its survival function

$$\bar{F}_{\text{GPD}}(x; \sigma, \xi) = \left(1 + \xi \frac{x}{\sigma}\right)_+^{-1/\xi}, \quad x \geq 0,$$

with $\sigma > 0$ and $(1 + \xi x)^{1/\xi} = e^x$ if $\xi = 0$ (Pickands, 1975). For $\xi < 0$, \bar{F}_{GPD} has support on $[0, \sigma/|\xi|]$. Condition (1), written as $X \in \text{MDA}_\xi$, means that X is in the maximum domain of attraction of an extreme value distribution with shape parameter ξ (see Resnick (1987)). In this case, the sequence of cumulative distribution functions of $a_u^{-1}(X - u) \mid X \geq u$ converges uniformly to $1 - \bar{F}_{\text{GPD}}$ on $[0, \infty)$. Thus, the distribution of exceedances above a large threshold u

30

(also called “peaks-over-threshold”) can be approximated in the following manner:

$$\text{pr}(X - u > x \mid X \geq u) = \text{pr}\{a_u^{-1}(X - u) > a_u^{-1}x \mid X \geq u\} \approx \bar{F}_{\text{GPD}}(x; \sigma a_u, \xi), \quad (2)$$

(Davison & Smith, 1990). This approximation, called the generalized Pareto approximation, is convenient in practice because it does not rely on a specific distributional assumption; X is only required to belong to some maximum domain of attraction, which holds for most common continuous distributions.

If the observations are discrete, however, one may want to preserve and utilize the discreteness in the extreme estimation. It is not, however, clear how discrete exceedances over high threshold should be modeled. The generalized Pareto approximation is often applied ignoring the discrete nature of the data. This poses two issues: first, a necessary condition for a discrete random variable X to be in some maximum domain of attraction in the case $x_F = \infty$ is that X is long-tailed,¹ i.e., $\bar{F}_X(u+1)/\bar{F}_X(u) \rightarrow 1$ as $u \rightarrow \infty$ (Shimura, 2012), and many common discrete distributions, including geometric, Poisson and negative binomial distributions, are not long-tailed. Specific convergence results for maxima of discrete observations have thus been derived (Anderson, 1970, 1980; Dkengne et al., 2016), but the limit is always a continuous distribution, which leads to the second issue: treating discrete data as continuous introduces a bias in the likelihood function. Since the shape and location parameters ξ and σ of the generalized Pareto approximation are unknown in practice, they must be estimated from the exceedance data. We will see that the bias may render the approximation inadequate — even when X is long-tailed, that is, when (2) is valid in theory.

Our contribution is to overcome these limitations by proposing two peaks-over-threshold methods, each relying on a parametric family of discrete distributions: the discrete generalized Pareto and the generalized Zipf distribution. The latter distributions exist in the literature but have not been justified for modeling extremes. As we will show, these new approximations can be theoretically motivated for X belonging to a broad class of discrete distributions, and they match or outperform the generalized Pareto approximation. They deliver similar results but it is still unclear if one of them should be preferred.

From now on, we assume that X is a discrete random variable with non-negative values, and $\xi \geq 0$. The first method adapts the condition $X \in \text{MDA}_\xi$ to the discrete case as follows. Suppose that there exists a random variable $Y \in \text{MDA}_\xi$ with survival function \bar{F}_Y on $[0, \infty)$ such that $\text{pr}(X \geq k) = \text{pr}(Y \geq k)$ for $k = 0, 1, 2, \dots$, that is, the equality in distribution, $X = \lfloor Y \rfloor$, holds. In this case, we say that X is in the discrete maximum domain of attraction, which we write as $X \in \text{D-MDA}_\xi$. We call Y an extension of X and such an extension is not unique. Shimura (2012) proved that $X \in \text{MDA}_\xi$ if and only if $X \in \text{D-MDA}_\xi$ and X is long-tailed.² It was also shown by Shimura (2012) that geometric, Poisson and negative binomial distributions belong to the discrete maximum domain of attraction. Therefore, $\text{MDA}_\xi \subsetneq \text{D-MDA}_\xi$ for discrete distributions. If $X \in \text{D-MDA}_\xi$ and $Y \in \text{MDA}_\xi$ is a corresponding extension satisfying $X = \lfloor Y \rfloor$ in distribution, then, for large integers u , we use (2) to obtain

$$\begin{aligned} \text{pr}(X - u = k \mid X \geq u) &= \text{pr}(Y - u \geq k \mid Y \geq u) - \text{pr}(Y - u \geq k + 1 \mid Y \geq u) \\ &\approx p_{\text{D-GPD}}(k; \sigma a_u, \xi), \end{aligned} \quad (3)$$

where $p_{\text{D-GPD}}$ is the probability mass function of the discrete generalized Pareto distribution defined by

$$p_{\text{D-GPD}}(k; \sigma, \xi) = \bar{F}_{\text{GPD}}(k; \sigma, \xi) - \bar{F}_{\text{GPD}}(k + 1; \sigma, \xi),$$

¹ All long-tailed distributions are heavy-tailed, but the converse is false.

² When X is long-tailed, an extension of X (which takes values in \mathbb{N}) can be X itself (if seen as taking values in \mathbb{R}).

for $k = 0, 1, 2, \dots$. Equation (3) provides a method for modeling discrete exceedances over threshold that we call the discrete generalized Pareto approximation. The latter distribution has been applied by Prieto et al. (2014) to model road accidents, while various aspects of discrete Pareto-type distributions were studied in Krishna & Pundir (2009), Buddana & Kozubowski (2014), and Kozubowski et al. (2015).

Whereas the first method is based on an extension of \bar{F}_X by a survival function in the maximum domain of attraction, the second method assumes instead an extension of p_X , the probability mass function of X . Suppose that there exists a non-negative random variable $Y \in \text{MDA}_{\xi/(1+\xi)}$ with survival function \bar{F}_Y on $[0, \infty)$ such that $p_X(k) = c \bar{F}_Y(k)$ for $k = d, d+1, d+2, \dots$, for some $c > 0$ and $d \in \mathbb{N}_0 = \{0, 1, \dots\}$. In this case, we say that p_X is in the discrete maximum domain of attraction which is denoted by $p_X \in \text{D-MDA}_{\xi/(1+\xi)}$, and call \bar{F}_Y an extension of p_X . We will show that $p_X \in \text{D-MDA}_{\xi/(1+\xi)}$ implies $X \in \text{MDA}_\xi$ (under a mild condition in the case $\xi = 0$), and that geometric, Poisson and negative binomial satisfy $p_X \in \text{D-MDA}_0$. It follows from (2) that, for large integers u ,

$$\text{pr}(X - u = k \mid X \geq u) = \frac{\text{pr}(Y > u + k) / \text{pr}(Y > u)}{\sum_{i=0}^{\infty} \text{pr}(Y > u + i) / \text{pr}(Y > u)} \approx p_{\text{GZD}}\{k; (1 + \xi)\sigma a_u, \xi\}, \quad (4)$$

where

$$p_{\text{GZD}}(k; \sigma, \xi) = \frac{(1 + \xi \frac{k}{\sigma})^{-1/\xi-1}}{\sum_{i=0}^{\infty} (1 + \xi \frac{i}{\sigma})^{-1/\xi-1}}, \quad k = 0, 1, 2, \dots, \quad (5)$$

is the probability function of a distribution that we call the generalized Zipf distribution. In the case $\xi = 0$, the latter is a geometric distribution (and so is the discrete generalized Pareto distribution), and in the case $\xi > 0$, it is a Zipf–Mandelbrot distribution (Mandelbrot, 1953). Zipf-type families have been fitted to various discrete datasets such as word frequencies (Booth, 1967), city sizes (Gabaix, 1999), company sizes (Axtell, 2001) and website visits (Clauset et al., 2009). The Zipf law, arising in the case $\xi = \sigma$, is sometimes presented as the discrete counterpart of the Pareto distribution (Arnold, 1983). We refer to the approximation procedure in (4) as the generalized Zipf approximation.

2. THEORETICAL RESULTS

We start by showing that the probability density and mass functions of the generalized Pareto, discrete generalized Pareto and Zipf distributions are asymptotically equivalent as σ tends to infinity. Proofs are given in the Supplementary Material.

PROPOSITION 1. For $\sigma > 0$, $\xi \geq 0$ and $q, \tilde{q} \in \{f_{\text{GPD}}, p_{\text{D-GPD}}, p_{\text{GZD}}\}$, it holds

$$\lim_{\sigma \rightarrow \infty} \sup_{k=0,1,2,\dots} \frac{q(k; \sigma, \xi)}{\tilde{q}(k; \sigma, \xi)} - 1 = 0.$$

This suggests that modeling a sample from $X - u \mid X \geq u$ by maximum likelihood using either f_{GPD} , $p_{\text{D-GPD}}$ or p_{GZD} should not differ too much if the estimated scale parameter $\hat{\sigma}$ is sufficiently large. When the sample size and u grow, $\hat{\sigma}$ only goes to infinity if the sequence a_u defined in (1) satisfies $a_u \rightarrow \infty$, which occurs if and only if X is long-tailed. Even in this case, a_u might grow too slowly for the three methods to be similar in practice, as we will see in Section 3.

The results below formally justify the approximation procedures we have introduced. We start with a convergence result for the discrete generalized Pareto approximation.

PROPOSITION 2. *If $X \in D\text{-MDA}_\xi$ for $\xi \geq 0$, then there exists a positive sequence $(a_u, u = 1, 2, \dots)$ such that*

$$\lim_{u \in \mathbb{N}_0, u \rightarrow \infty} \sup_{k=0,1,2,\dots} pr(X = u + k \mid X \geq u) - p_{D\text{-GPD}}(k; a_u, \xi) = 0. \quad (6)$$

We remark that (6) is not informative if $a_u \rightarrow \infty$ because the two terms converge to 0. Next, we consider the case $p_X \in D\text{-MDA}$. Recall that a distribution F is in MDA_0 if and only if the survival function has a representation

$$\bar{F}(x) = c(x) \exp \left\{ - \int_0^x \frac{1}{a(y)} dy \right\}, \quad x \in \mathbb{R}, \quad (7)$$

where $a(\cdot)$, called the auxiliary function, is positive and differentiable with $a'(x) \rightarrow 0$ as $x \rightarrow \infty$; and $c(\cdot)$ is a positive function with limit $c > 0$ (Embrechts et al., 2013). If $c(x) = c$ on (d, ∞) for some $d \in \mathbb{R}$, then we say that the distribution F satisfies the von Mises condition.

THEOREM 1. *If $p_X \in D\text{-MDA}_{\xi/(1+\xi)}$ and $\xi > 0$, then $X \in \text{MDA}_\xi$ and, for any sequence of nonnegative integers $(k_u)_{u \in \mathbb{N}_0}$ such that $\sup_u k_u/u < \infty$,*

$$\lim_{u \in \mathbb{N}, u \rightarrow \infty} \frac{pr(X = k_u + u \mid X \geq u)}{q(k_u; \xi u, \xi)} = 1. \quad (8)$$

where $q \equiv f_{GPD}, p_{D\text{-GPD}}$ and p_{GZD} .

If $p_X \in D\text{-MDA}_0$ and if the auxiliary function of an extension \bar{F} of p_X satisfies $\lim_{x \rightarrow \infty} a(x) = \sigma > 0$, then $X \in D\text{-MDA}_0$ and

$$\lim_{u \in \mathbb{N}, u \rightarrow \infty} pr(X = k + u \mid X \geq u) = p_{D\text{-GPD}}(k; \sigma, 0) = p_{GZD}(k; \sigma, 0), \quad k = 0, 1, 2, \dots \quad (9)$$

The condition $p_X \in D\text{-MDA}$ is satisfied, among others, by the Zipf–Mandelbrot, geometric, Poisson and negative binomial distributions as shown below and in the Supplementary Material.

Example 1. The probability mass function of a Zipf–Mandelbrot distribution is proportional to $(k + q)^{-1-1/\xi}$ for $k = 0, 1, 2, \dots, q > 0, \xi > 0$, and satisfies $p_X \in D\text{-MDA}_{\xi/(1+\xi)}$ because it can be extended by $\bar{F}_Y(y) = c(y + q)^{-1-1/\xi}$ for $y \geq 0$ and some $c > 0$ with $Y \in \text{MDA}_{\xi/(1+\xi)}$. The probability mass function of a geometric distribution belongs to $D\text{-MDA}_0$ as it coincides up to a constant with the survival function of an exponential distribution. The latter distribution clearly satisfies the von Mises condition and thus is a member of MDA_0 . The auxiliary function is, in fact, equal (eventually) to $1/\lambda$, where λ is the rate of the exponential distribution.

To summarize, for a discrete random variable X and $\xi \geq 0$, it holds $X \in \text{MDA}_\xi$ if and only if $X \in D\text{-MDA}_\xi$ and X is long-tailed. If $\xi > 0$, then $p_X \in D\text{-MDA}_{\xi/(1+\xi)}$ implies $X \in D\text{-MDA}_\xi$; the same implication holds in the case $\xi = 0$ if the auxiliary function of the extension of p_X satisfies $a(x) \rightarrow \sigma \in (0, \infty)$ as $x \rightarrow \infty$.

3. EMPIRICAL RESULTS

We assess the performance of the discrete generalized Pareto and the generalized Zipf approximations for estimating the probability of a rare event from discrete data, and illustrate why they should be preferred to the generalized Pareto approximation, whether X is long-tailed or not. Let $\alpha = 2, \beta = 0.75$ and

$$X = \lfloor Y \rfloor, \quad Y \sim \text{Inverse-gamma}(\alpha, \beta), \quad (10)$$

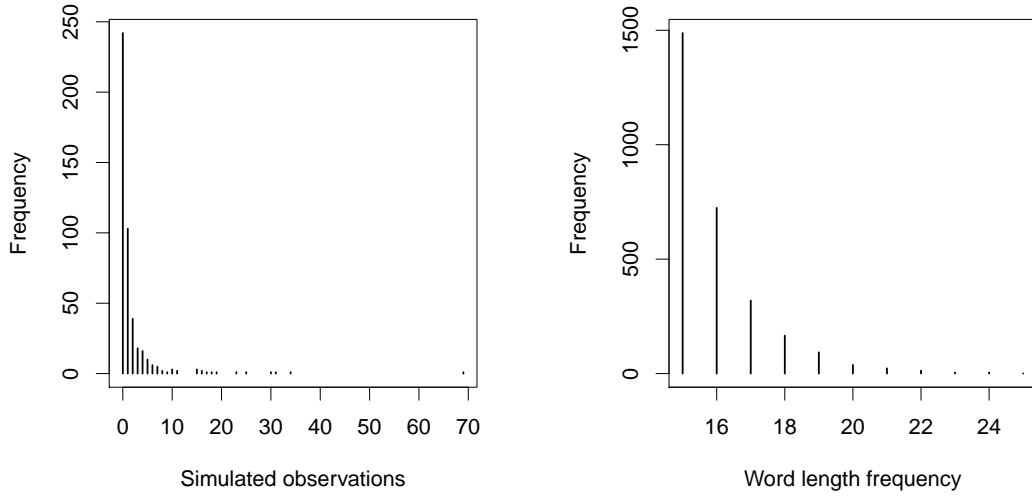


Fig. 1: On the left: frequency plot of 462 exceedances of X over the 95th empirical quantile $u = 2$ simulated from (10). On the right: frequency plot of the length of the 2875 longest French words.

where the inverse-gamma distribution has density function $f(x) = \Gamma(\alpha)^{-1} \beta^\alpha x^{-\alpha-1} \exp(-\beta/x)$, $x > 0$. The experiment described below is repeated 500 times. An independent and identically distributed sample of size 8000 is drawn from the distribution of X . The goal is to estimate the probability of the extreme region

$$p_e = \text{pr}(X \geq \lfloor q_e \rfloor), \quad \lfloor q_e \rfloor = 52, \quad (11)$$

where q_e is the 99.99 percentile of Y , i.e., the value exceeded once every 10 000 times on average. The strategy pursued is to select an integer threshold u as the 95th empirical percentile of the sample,³ fit parametric distributions to the exceedances $X - u \mid X \geq u$, and use them to extrapolate the tail and estimate p_e . It clearly holds $p_X \in \text{D-MDA}_{\xi/(1+\xi)}$ for $\xi = 1/\alpha = 1/2$, thus the three approximations are justified. The generalized Pareto distribution is fitted to the observations shifted by continuity correction $\delta = 0$ or $\delta = \frac{1}{2}$. As a benchmark, we will also estimate p_e from a sample of the continuous variable Y (as opposed to its discretization X) using the generalized Pareto approximation.

A frequency plot of the exceedances of a sample of X above u is displayed on the left-hand side in Figure 1. For each model, we compute the maximum likelihood estimators $\hat{\sigma}$ and $\hat{\xi}$ by performing a two dimensional maximization using the function `optim` of R (R Core Team, 2015) with starting values $(1, 1)$. We then compute \hat{p}_e and approximate 90% confidence intervals from the Fisher information matrix under asymptotic normality of the estimators. Table 1 displays: the average estimates \hat{p}_e , $\hat{\xi}$ and $\hat{\sigma}$ over the 500 replications of the experiment, the coverage⁴ of the confidence intervals, their average length and their true length.⁵ The discrete generalized Pareto and Zipf approximations provide relatively accurate estimates of p_e from the discretized data

³ Selecting an appropriate threshold is crucial when estimating high quantiles and can be based on techniques such as mean residual plots (see e.g. Davison & Smith (1990)).

⁴ Coverage indicates the proportion of time the truth lies in the confidence interval.

⁵ True length is here defined as $\ell^* = q_{0.05}(\hat{\mathbf{p}}_e) - q_{0.95}(\hat{\mathbf{p}}_e)$, where $\hat{\mathbf{p}}_e$ is the vector of maximum likelihood estimates in the 500 replicated experiments, and $q_{(\cdot)}$ is the quantile function.

Table 1: Performance of several methods in estimating the probability p_e of the rare event defined in (11) from about 460 exceedances in each experiment. The table displays average maximum likelihood estimators for p_e , ξ and σ across the 500 replicated experiments. Coverage c , average length l and true length l^* of 90% confidence intervals are shown between brackets. The discrete generalized Pareto and Zipf approximations are superior in this case.

Truth	$\hat{p}_e \cdot 10^4 (c, l, l^*)$ 1.03	$\hat{\xi} (c, l)$ 0.50	$\hat{\sigma} (l)$
Fitted to $Y - u \mid Y \geq u$			
Generalized Pareto distribution	1.07 (85%, 1.84, 1.81)	0.49 (94%, 0.28)	1.14 (0.36)
Fitted to $X - u \mid X \geq u, X = \lfloor Y \rfloor$			
Discrete generalized Pareto distribution	1.09 (87%, 1.92, 1.85)	0.49 (93%, 0.29)	1.13 (0.40)
Generalized Zipf distribution	1.11 (88%, 1.97, 1.88)	0.50 (94%, 0.30)	1.34 (0.35)
Generalized Pareto distribution, $\delta = \frac{1}{2}$	0.44 (31%, 0.86, 0.97)	0.36 (35%, 0.22)	1.38 (0.39)
Generalized Pareto distribution, $\delta = 0$	50.42 (85%, 162.32, 71.45)	8.29 (0%, 1.58)	0.00 (0.00)

Table 2: Fit of several distributions to the length of the 2875 longest French words, and to the number of extreme tornadoes per outbreak for the 435 outbreaks with 12 or more such tornadoes in the United States between 1965 and 2015. The table displays p-value of discrete Kolmogorov–Smirnov tests (of the discretized model in the case of continuous models), negative log-likelihood $-\ell$ and maximum likelihood estimates with 90% confidence intervals and possible temporal trend $\hat{\sigma}_t$ in the scale parameter.

	p-val.	$-\ell$	ξ	$\hat{\sigma}_0$	$\hat{\sigma}_t$
Word length					
Discrete generalized Pareto dist.	0.40	3894.0	0.02 _[-0.01,0.06]	1.36 _[1.30,1.43]	
Generalized Zipf distribution	0.40	3894.0	0.02 _[-0.01,0.06]	1.37 _[1.32,1.43]	
Generalized Pareto dist., $\delta = \frac{1}{2}$	0.02	3951.2	-0.04 _[-0.06,-0.01]	1.51 _[1.45,1.57]	
Negative binomial	0.37	3893.9			
Tornado outbreak					
Discrete generalized Pareto dist.	0.19	1439.92	0.27 _[0.16,0.37]	4.81 _[3.64,5.99]	6.11 _[3.74,8.48]
Generalized Pareto dist., $\delta = \frac{1}{2}$	0.18	1439.93	0.26 _[0.16,0.37]	4.86 _[3.68,6.04]	6.13 _[3.75,8.50]

with a coverage close to the correct one of 90%, and their performance is good relative to the situation of full information where the continuous data are available — notice how the estimates are very similar to one another. On the other hand, the two versions of the generalized Pareto approximation perform poorly, the worst being the case $\delta = 0$.

The ability of the discrete generalized Pareto and Zipf approximations to accurately estimate the probability of rare events is supported by complementary simulated cases covering $\xi = 0$ and $\xi < 0$ (Hitz, 2016, Chapter 2), and illustrated here on three real datasets. The first consists of the frequency X of word length in the French lexicon (New et al., 2004); for instance, “anti-constitutionnellement” is the only word of 25 letters in French. We focus on describing the tail distribution and fit the usual models to $X - u \mid X \geq u$ with $u = 15$, the 98th percentile of the data. A frequency plot of the 2875 exceedances is shown on the right-hand side of Figure 1. The discrete generalized Pareto and Zipf distributions deliver a good fit and similar estimations to one another, and clearly outperform the generalized Pareto approximation as shown in Table 2 by p-values⁶ of discrete Kolmogorov–Smirnov tests based on the difference between the fitted

⁶ Procedure for computing p-values in Table 2 for the word length data: resample the data with replacement; compute the difference between the fitted and empirical distribution of this sample; get a p-value by Monte Carlo simulation using R package `dgoF` (Arnold & Emerson, 2011); repeat 200 times and take the average.

Table 3: On the left: frequency table of multiple births in the United States from 1995 to 2014. On the right: performance of several methods in estimating the probability p_e of an American women delivering quintuplets or more at birth using only one thousandth of the dataset. In each experiment, the threshold was $u = 1$ and there were about 2600 exceedances (see Table 1 for notation). The discrete generalized Pareto and Zipf provide useful techniques for such extrapolations.

	Multiple Births		$\hat{p}_e \cdot 10^5 (c, l, l^*)$
Single	78 178 588	Truth	1.7
Twin	2 500 340	Discrete generalized Pareto distribution	1.4 (74%, 2.9, 2.9)
Triplet	1 17 603	Generalized Zipf distribution	1.6 (87%, 3.3, 2.8); n/a ⁹
Quadruplet	8 108	Negative Binomial	1.2 (65%, 2.3, 2.3)
Quint. or more	1 353	Generalized Pareto distribution, $\delta = \frac{1}{2}$	n/a ¹⁰

distribution and the empirical distribution of bootstrapped data. Notice that the negative binomial also fits well in this case.

The second dataset comes from Tippett et al. (2016) who report the number X of extreme tornadoes per outbreak in the United States between 1965 and 2015, where an outbreak is a sequence of tornadoes that are high on the Fujita scale and occur close to each other in time. The authors found that the 435 observations from $X - u \mid X \geq u$ for $u = 12$ were well modeled by a generalized Pareto distribution with linear temporal trend in the scale parameter⁷ and continuity correction $\delta = \frac{1}{2}$. Maximum likelihood estimates and discrete Kolmogorov–Smirnov tests⁸ in Table 2 show that there is virtually no difference between the three approximations (only two of them are presented). This is consistent with Proposition 1 since the location parameter $\hat{\sigma}$ is larger in this case. Treating the tornado data as continuous is acceptable because there are fewer tied observations: about 38% of the data consists of values shared with no more than 20 other observations, compared to 13% for the simulated data and 1% for the word length data. Loosely, the data look less discrete (a frequency plot is displayed in the Supplementary Material), thus the generalized Pareto approximation is appropriate here.

The third dataset counts the number X of multiple births in the United States from 1995 to 2014 and is displayed on the left-hand side of Table 3 (Hamilton et al., 2015). The observations are censored from above and only take 5 distinct values, it is thus interesting to see if the discrete generalized Pareto and Zipf distributions can still describe the tail of the data in this non-standard estimation problem. We randomly select from the dataset a sample that contains a thousand times fewer observations, and estimate from these the probability p_e that an American women delivers quintuplets or more by fitting a right-censored version of the usual models to $X^C - u \mid X^C \geq u$ for $u = 2$, where $X^C = \min(X, 5)$. The experiment was repeated 500 times, and each sample contained on average 9 quatruplets and 1 quintuplet or more. Table 3 shows that the discrete generalized Pareto and Zipf distributions outperform common alternatives, and seem to be useful techniques for inference from such limited data. The applicability of peaks-over-threshold methods when u is a particularly small integer should be more rigorously studied.

Future work could explore the use of the generalized Zipf and discrete Pareto distributions in the case $\xi < 0$, and further investigate how they relate to each other as they seem to perform similarly. The latter distribution benefits from its closed-form survival and probability mass function,

⁷ The scale parameter is modeled as $\sigma(\mathbf{t}) = \sigma_0 + \sigma_1 \mathbf{t}$, where \mathbf{t} is the time covariate rescaled between $[0, 1]$.

⁸ Procedure for computing p-values in Table 2 for the tornado data: split the dataset into 5 groups depending on which time covariates are the nearest to $t_i = 0.1, 0.3, 0.5, 0.7, 0.9$; for each group, assume $\sigma(\mathbf{t}) = \sigma_0 + \sigma_1 t_i$ and compute the p-value of a discrete Kolmogorov–Smirnov test as explained previously; report the smallest of these 5 p-values.

⁹ In 52 out of 500 replicated experiments, maximum likelihood estimates could not be computed numerically. In 70 experiments, the hessian matrix could not be computed numerically.

¹⁰ The log-likelihood function could not be maximized numerically.

allowing for exact likelihood based inference. In conclusion, there is no downside to fit a discrete generalized Pareto for discrete data as opposed to a generalized Pareto distribution.

ACKNOWLEDGEMENT

The first author is grateful to the Berrow Foundation for financial support and would like to thank Robin Evans, Gesine Reinert, David Steinsaltz and Jonathan Tawn for valuable advice and encouragement. We thank as well Michael Tippett for providing the tornado data. This research was partially supported by the ARO grant W911NF-12-10385.

REFERENCES

- ANDERSON, C. W. (1970). Extreme value theory for a class of discrete distributions with applications to some stochastic processes. *Journal of Applied Probability* **7**, 99–113.
- ANDERSON, C. W. (1980). Local limit theorems for the maxima of discrete random variables. *Mathematical Proceedings of the Cambridge Philosophical Society* **88**, 161–165.
- ARNOLD, B. C. (1983). *Pareto distribution*. International Cooperative Publishing House, Maryland.
- ARNOLD, T. B. & EMERSON, J. W. (2011). Nonparametric goodness-of-fit tests for discrete null distributions. *The R Journal* **3**, 34–39.
- AXTELL, R. L. (2001). Zipf distribution of US firm sizes. *Science* **293**, 1818–1820.
- BOOTH, A. D. (1967). A law of occurrences for words of low frequency. *Information and Control* **10**, 386–393.
- BUDDANA, A. & KOZUBOWSKI, T. J. (2014). Discrete Pareto distributions. *Economic Quality Control* **29**, 143–156.
- CLAUSET, A., SHALIZI, C. R. & NEWMAN, M. E. J. (2009). Power-law distributions in empirical data. *SIAM Review* **51**, 661–703.
- DAVISON, A. C. & SMITH, R. L. (1990). Models for exceedances over high thresholds. *Journal of the Royal Statistical Society. Series B. Methodological* **52**, 393–442.
- DKENGNE, P. S., ECKERT, N. & NAVEAU, P. (2016). A limiting distribution for maxima of discrete stationary triangular arrays with an application to risk due to avalanches. *Extremes* **19**, 25–40.
- EMBRECHTS, P., KLÜPPELBERG, C. & MIKOSCH, T. (2013). *Modelling extremal events: for insurance and finance*. Springer, New York.
- GABAIX, X. (1999). Zipf’s law and the growth of cities. *The American Economic Review* **89**, 129–132.
- HAMILTON, B. E., MARTIN, J. A., OSTERMAN, M. J., CURTIN, S. C. & TJ, M. (2015). Births: Final data for 2014. *National Vital Statistics Reports* **64**.
- HITZ, A. S. (2016). *Modelling of Extremes*. Ph.D. thesis, University of Oxford.
- KOZUBOWSKI, T. J., PANORSKA, A. K. & FORISTER, M. L. (2015). A discrete truncated Pareto distribution. *Statistical Methodology* **26**, 135–150.
- KRISHNA, H. & PUNDIR, P. S. (2009). Discrete Burr and discrete Pareto distributions. *Statistical Methodology* **6**, 177–188.
- MANDELBROT, B. (1953). Contribution à la théorie mathématique des jeux de communication. *Publications de l’Institut de statistique de l’Université de Paris*.
- NEW, B., PALLIER, C., BRYBAERT, M. & FERRAND, L. (2004). Lexique 2: A new French lexical database. *Behavior Research Methods, Instruments, & Computers* **36**, 516–524.
- PICKANDS, III, J. (1975). Statistical inference using extreme order statistics. *The Annals of Statistics* **3**, 119–131.
- PRIETO, F., GÓMEZ-DÉNIZ, E. & SARABIA, J. M. (2014). Modelling road accident blackspots data with the discrete generalized Pareto distribution. *Accident Analysis & Prevention* **71**, 38–49.
- R CORE TEAM (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing.
- RESNICK, S. I. (1987). *Extreme values, regular variation, and point processes*. Springer, New York.
- SHIMURA, T. (2012). Discretization of distributions in the maximum domain of attraction. *Extremes* **15**, 299–317.
- TIPPETT, M. K., LEPORE, C. & COHEN, J. E. (2016). More tornadoes in the most extreme US tornado outbreaks. *Science* **354**, 1419–1423.