

An extended generalized Pareto regression model for count data

Touqeer Ahmad¹, Carlo Gaetan² and Philippe Naveau³

¹ ENSAI, CREST-UMR9194, Bruz, France.

² Dipartimento di Scienze Ambientali, Informatica e Statistica, Università Ca' Foscari-Venezia, Italy.

³ Laboratoire des Sciences du Climat et de l'Environnement, UMR 8212, CEA-CNRS-UVSQ, IPSL & Université Paris-Saclay, Gif-sur-Yvette, France.

Address for correspondence: ENSAI, 51 Rue Blaise Pascal, 35170, Bruz, France.
E-mail: touqeer.ahmad@ensai.fr.

Abstract: The statistical modeling of discrete extremes has received less attention than their continuous counterparts in the Extreme Value Theory (EVT) literature. One approach to the transition from continuous to discrete extremes is the modeling of threshold exceedances of integer random variables by the discrete version of the generalized Pareto distribution. However, the optimal choice of thresholds defining exceedances remains a problematic issue. Moreover, in a regression framework, the treatment of the majority of non-extreme data below the selected threshold is either ignored or separated from the extremes. To tackle these issues, we expand on the concept of employing a smooth transition between the bulk and the upper tail of the distribution. In the case of zero inflation, we also develop models with an additional parameter. To incorporate possible predictors, we relate the parameters to additive smoothed predictors via an appropriate link, as in the generalized additive model (GAM) framework. A penalized maximum likelihood estimation procedure is implemented. We illustrate our modeling proposal with a real dataset of avalanche activity in the French Alps. With the advantage of bypassing the threshold selection step, our results indicate that the proposed models are more flexible and robust than competing models, such as the negative binomial distribution.

Key words: Extreme value theory; discrete extended generalized Pareto distribution; zero-inflated models; generalized additive models

1 Introduction

Modeling count data, i.e., non-negative integers, in the presence of covariates is a very common task in many research areas. The Poisson regression model is often of limited use because count data typically exhibit overdispersion (i.e., the variance of the counts appears larger than the mean) and/or an excessive number of zeros. Additional parameters can be inserted to deal with overdispersion, e.g., the quasi-Poisson model (Wedderburn, 1974), or different distributions, such as negative binomial distribution, can be fitted. To model zero inflation, Lambert (1992) studied a two-component mixture model, one component is a point mass at zero and the other component is an assumed parametric count distribution. Lambert’s specification is an example of a distributional regression model (Stasinopoulos et al., 2018; Kneib et al., 2023). The term “distributional” emphasizes that several characteristics of the conditional distribution of the data are modeled in terms of covariates, rather than only the mean.

So far, the main focus of the literature has been on the relationship between the location, the scale, and the shape with the covariates (Rigby and Stasinopoulos, 2005), less attention has been paid to the tail of the count distribution.




Figure 1: (a) Frequency table plot (\log_{10} scale) of avalanche events in the Haute-Maurienne massif of the French Alps, (b) Q-Q plot of randomized residuals from the zero-inflated negative binomial model with additive covariates. The dotted lines show the 95% point-wise confidence intervals.

As a motivating example, we consider a dataset on the avalanche activity in the Haute-Maurienne region of the French Alps. The avalanche activity is measured by a three-day moving sum of the avalanche events recorded from February 1982 to April 2021, see Evin et al. (2021) for details. Figure 1(a) shows a large relative frequency

of zeros (meaning no avalanche has been reported) as well as heavy-tailed behavior.

We fit a zero-inflated negative binomial regression model under the framework of generalized additive models for location, scale, and shape (GAMLSS) where the parameters are related to additive environmental covariates (see Section 4 for a detailed description) via suitable link functions (Stasinopoulos et al., 2018). The randomized quantile residuals (Dunn and Smyth, 1996) are used to check the adequacy of the fitted model. Figure 1(b) clearly shows that the fitted models do not correctly estimate the upper tail behavior of avalanche extremes. In addition, the number of zeros is not correctly predicted in this example.

Extreme value theory, originally developed by Fisher and Tippett (1928), provides a mathematical blueprint to model very high and very low-frequency events (e.g., extreme temperatures, heavy rainfall intensities, heavy floods, and extreme winds, etc.), and monographs such as Coles (2001) or Beirlant et al. (2004) discuss the main extreme value models. In particular, under the *peak-over-threshold* (POT) approach (Pickands, 1975), the distribution of exceedances of a high threshold is often approximated by the Generalized Pareto Distribution (GPD). Modifications of GPD to discrete data exist in the literature (Krishna and Pundir, 2009; Buddana and Kozubowski, 2014; Kozubowski et al., 2015), and recently Hitz et al. (2024) discussed discrete versions of GPD to approximate the tail behavior of integer-valued random variables. This approach still requires the definition of a threshold at a high quantile, which is not easy due to the discrete nature of the data (Daouia et al., 2023).

It should also be noted that especially environmental time series are rarely stationary and depend on environmental factors. A standard approach to modeling continuous extremes of a non-stationary process focuses on maintaining a predetermined threshold but treating parameters of the GPD as functions of covariates (Davison and Smith, 1990). An alternative approach (Eastoe and Tawn, 2009) uses preprocessing methods to model the non-stationarity in the body of the process to produce transformed data and then uses standard methods to model the extremes of the transformed data. The first approach has been adapted to the discrete case by Ranjbar et al. (2022). The second approach seems to be difficult to adapt. The distribution of the preprocessed data cannot be connected to a distribution of count data.

The proposed model addresses the issue of the POT approach ignoring or separating non-extreme data below the selected threshold from the extremes. The model utilizes a smooth transition between the bulk and upper tail of the distribution, for the full range of the data, while bypassing a threshold selection. The discrete extended version of GPD (DEGPD) is derived by discretizing the cumulative distribution function (CDF) of an extended GPD (Naveau et al., 2016). The model takes into account the possible effects of covariates in a non-parametric way. Since it is possible to have a dataset with an excess of zeros, such as in the motivating example, we also consider a mixture of the previous distribution with a degenerate distribution at zero. This results in a distribution named Zero-Inflated DEGPD (ZIDEGPD).

The paper is organized as follows. Section 2 introduces the DEGPD. The extension to deal with many zeros and covariate effects is given in Section 3. Section 4 discusses applications of DEGPD and ZIDEGPD to avalanche data with environmental covariates. Finally, Section 5 concludes with a summary of our results and a discussion of future research directions.

2 Discrete extreme modelling

The distribution of exceedances (i.e., the amount of data that appears over a given high threshold) is often approximated by the Generalized Pareto distribution (GPD) defined by its CDF as

$$F(z; \sigma, \xi) = \begin{cases} 1 - (1 + \xi z/\sigma)_+^{-1/\xi} & \xi \neq 0 \\ 1 - \exp(-z/\sigma) & \xi = 0 \end{cases}, \quad (2.1)$$

where $(a)_+ = \max(a, 0)$. The $\sigma > 0$ and $-\infty < \xi < +\infty$ represent the scale and shape parameters of the distribution, respectively.

More precisely let X be a random variable taking values in $[0, x_F)$ where $x_F \in (0, \infty) \cup \{\infty\}$. Suppose that there exists a strictly positive sequence a_u such that the distribution of $a_u^{-1}(X - u)|X \geq u$ weakly converge to a non-degenerate probability distribution on $[0, \infty)$ as $u \rightarrow x_F$, then this distribution is the GPD (Balkema and De Haan, 1974). Thus, for large u ,

$$\Pr(X - u > x|X \geq u) = \Pr(a_u^{-1}(X - u) > a_u^{-1}x|X \geq u) \approx 1 - F(x; a_u\sigma, \xi).$$

The shape parameter, ξ , defines the tail behavior of the GPD. If $\xi < 0$, the upper tail is bounded. If $\xi = 0$, we have the exponential distribution, where all moments are finite. If $\xi > 0$, the upper tail is unbounded and the higher moments ultimately become infinite. The three defined cases are labeled “short-tailed”, “light-tailed”, and “heavy-tailed”, respectively. These categorizations enhance the flexibility of the GPD and underscore its adaptability to various modeling scenarios.

Using the GPD to approximate the distribution tail for discrete data can be inappropriate, as pointed out in Hitz et al. (2024). These authors proposed to approximate the distribution tail of a count random variable Y by discretizing the CDF defined by (2.1) and, for large u ,

$$\Pr(Y - u = k|Y \geq u) = F(k + 1; \sigma, \xi) - F(k; \sigma, \xi), \quad k \in \mathbb{N}_0, \quad (2.2)$$

with $\sigma > 0$ and $\xi \geq 0$. The distribution is called discrete GPD (DGPD), and several properties of discrete Pareto type distributions have been studied previously in the literature (Krishna and Pundir, 2009; Buddana and Kozubowski, 2014; Kozubowski et al., 2015).

A drawback of GPD in the continuous case is that it only models observations that occur above a certain high threshold. This imposes an artificial dichotomy in the data (i.e., observations are either below or above the threshold), and finding the optimal threshold remains complex for practitioners. The choice becomes more complicated when the observations feature a substantial number of ties.

In the continuous extreme value setting, many authors have attempted to model the full range of data without threshold selection (Frigessi et al., 2002; Carreau and Bengio, 2009; MacDonald et al., 2011; Papastathopoulos and Tawn, 2013; Naveau et al., 2016; Stein, 2021).

Notably, Papastathopoulos and Tawn (2013) proposed an extension of GPD that incorporated an additional shape parameter without affecting the tail behavior. The inclusion of this parameter stabilized the GPD parameter estimates for threshold selection, allowing a lower threshold to be selected. In order to work with the modeling of the lower tail and the bulk of the distribution, Naveau et al. (2016) identified two conditions that ensured compliance with the EVT for the lower and upper tails. The transition between the two tails was modeled by a function on $[0, 1]$, which can take different forms. For example, Tencaliec et al. (2020) worked with a Bernstein polynomial base. de Carvalho et al. (2022) extended this type of approach to Bayesian lasso structures. The main ingredient of all these constructions, called extended generalized Pareto distributions (EGPDs), is the idea of the integral transformation to simulate GPD random draws, i.e. $F_{\sigma,\xi}^{-1}(U)$, where $U \sim \mathcal{U}(0, 1)$ represents a uniformly distributed random variable on $(0, 1)$ and $F_{\sigma,\xi}^{-1}$ denotes the inverse of the CDF (2.1). This leads to the family of distribution for the random variable

$$Z = F_{\sigma,\xi}^{-1}(G^{-1}(U)), \tag{2.3}$$

where G is a CDF on $[0, 1]$ and $U \sim \mathcal{U}(0, 1)$. Clearly, the CDF of Z is $G(F(z; \sigma, \xi))$. The key problem is to find a function G which preserves the upper tail behavior with shape parameter ξ and also controls the lower tail behavior. Naveau et al. (2016) defined restrictions for validity of G functions. For instance, the tail of G denoted by $\bar{G} = 1 - G$ has to satisfy

$$\lim_{u \rightarrow 0} \frac{\bar{G}(1 - u)}{u} = a, \text{ for some finite } a > 0 \text{ (upper tail behavior),} \tag{2.4}$$

$$\lim_{u \rightarrow 0} \frac{G(u)}{u^c} = c, \text{ for some finite } c > 0 \text{ (lower tail behavior).} \tag{2.5}$$

Four parametric examples for G have been proposed in Naveau et al. (2016) (see also below). We follow the same idea and define the probability mass function (pmf) for the count variable Y as

$$\Pr(Y = y) = G(F(y + 1; \sigma, \xi)) - G(F(y; \sigma, \xi)), \quad y \in \mathbb{N}_0. \tag{2.6}$$

The distribution defined by (2.6) is referred to as the discrete extended generalized Pareto distribution (DEGPD). The explicit formula of CDF of DEGPD is developed

as

$$\Pr(Y \leq y) = G(F(y + 1; \sigma, \xi)) \quad (2.7)$$

and the quantile function is derived as

$$q_p = \begin{cases} \left\lceil \frac{\sigma}{\xi} \left\{ (1 - G^{-1}(p))^{-\xi} - 1 \right\} \right\rceil - 1, & \text{if } \xi > 0 \\ \lceil -\sigma \log(1 - G^{-1}(p)) \rceil - 1, & \text{if } \xi = 0 \end{cases} \quad (2.8)$$

with $0 < p < 1$. In this paper, we use four parametric expressions of $G(\cdot)$ (Naveau et al., 2016), namely

Model (i): $G(u; \psi) = u^\kappa$, $\psi = \kappa > 0$;

Model (ii): $G(u; \psi) = 1 - D_\delta\{(1 - u)^\delta\}$, $\psi = \delta > 0$ where D_δ is the CDF of a Beta random variable with parameters $1/\delta$ and 2, that is:

$$D_\delta(u) = \frac{1 + \delta}{\delta} u^{1/\delta} \left(1 - \frac{u}{1 + \delta}\right);$$

Model (iii): $G(u; \psi) = [1 - D_\delta\{(1 - u)^\delta\}]^{\kappa/2}$, $\psi = (\delta, \kappa)$ with $\delta > 0$ and $\kappa > 0$;

Model (iv): $G(u; \psi) = pu^{\kappa_1} + (1 - p)u^{\kappa_2}$, $\psi = (p, \kappa_1, \kappa_2)$ with $\kappa_2 \geq \kappa_1 > 0$ and $p \in (0, 1)$.

The parametric family (i) leads to a pmf of DEGPD with three parameters (κ, σ and ξ): κ deals the shape of the lower tail, σ is a scale parameter, and ξ controls the rate of upper tail decay. Thus, Figure 2(a) shows the behavior of the pmf of DEGPD with fixed scale and upper tail shape parameter (i.e., $\sigma = 1$ and $\xi = 0.5$) and with different values of lower tail behaviors ($\kappa=1, 2, 5$). The DGPD is recovered when $\kappa = 1$, and additional flexibility for low values is attained by varying κ . For instance, more flexibility on the lower tail can be observed without losing upper tail behavior in Figure 2(a) when $\kappa = 5$.

The parametric family (ii) is another interesting choice for constructing DEGPDs. This choice is more complex than the previous two. Figure 2(b) illustrates the behavior of pmf with different values of δ . In the continuous setting, we observe that the EGPD associated with this G family converges to GPD as δ increases to infinity. Furthermore, the conditions (2.4) and (2.5) are also satisfied with $\delta = 2$, see Naveau et al. (2016) for more details. In discrete settings, the DEGPD corresponding $G(u; \psi) = 1 - D_\delta\{(1 - u)^\delta\}$ also become very close to the pmf of DGPD when δ increases to infinity. In general, the parameter δ describes the central part of the distribution. Thus, this parameter relatively improves the flexibility for the central part of the distribution. The δ parameter can be interpreted as a threshold tuning parameter. One of the drawbacks of DEGPD (ii) is that it models only the central

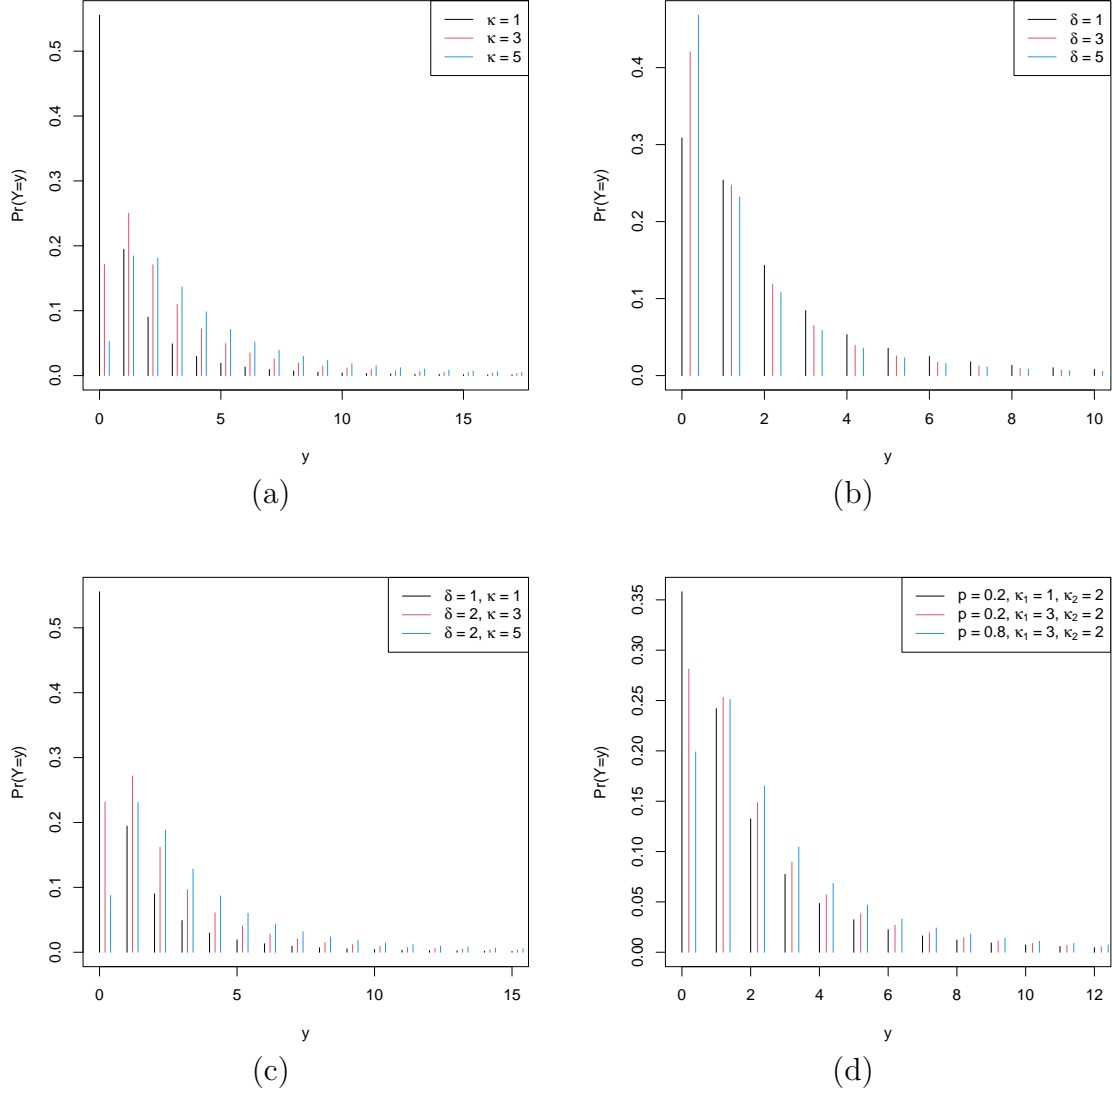


Figure 2: Probability mass function corresponding to model (2.6) for $\sigma = 1$, $\xi = 0.5$. (a) Model (i): $G(u; \psi) = u^\kappa$, $\kappa = 1, 3, 5$; (b) Model (ii): $G(u; \psi) = 1 - D_\delta(1 - u)^\delta$ $\delta = 1, 3, 5$; (c) Model (iii): $G(u; \psi) = [1 - D_\delta(1 - u)^\delta]^{\kappa/2}$ for $\delta = 1, 2, 2$ and $\kappa = 1, 3, 5$, and (d) Model (iv): $G(u; \psi) = pu^{\kappa_1} + (1 - p)u^{\kappa_2}$ for $p = 0.2, 0.2, 0.8$, $\kappa_1 = 1, 3, 3$ and $\kappa_2 = 2, 2, 2$.

and upper parts of the distribution. On the other hand, the behavior of the lower tail could not be estimated directly.

The parametric family (iii) supports the lower tail of the distribution with the $\kappa > 0$ parameter. Interestingly, this family also tends to the DEGPD with parameters $(\kappa, \delta, \sigma$ and $\xi)$. The $(\kappa, \delta$ and $\xi)$ represents the lower, central, and upper parts of the distribution, respectively, and σ is a scale parameter as usual. In particular, Figure 2(c) shows the behavior of pmf of DEGPD linked with $G(u; \psi) = [1 - D_\delta\{(1 - u)^\delta\}]^{\kappa/2}$ at different settings of the parameters. Changing the values of κ also reveals flexibility of the lower tail.

The parametric family (iv) is the mixture of power laws: κ_1 identifies the lower tail shape, κ_2 modifies the central distribution shape, and σ and ξ are scale and upper tail parameters, respectively. It can be observed from Figure 2(d) that the DEGPD related to $G(u; \psi) = pu^{\kappa_1} + (1 - p)u^{\kappa_2}$ is also showing flexibility with $p = 0.2, 0.8$, $\sigma = 1$, $\xi = 0.5$, $\kappa_1 = 1, 3$, constant values of $\kappa_2 = 2$.

3 Zero-inflation and regression modeling

As we have seen for our motivating example, many zeros can be found in various real data sets. In that case, the current model with a flexible lower and upper tail cannot be adjusted for the excessive zeros. We follow Lambert (1992) and change DEGPD's pmf to

$$\Pr(Y = y) = \begin{cases} \pi + (1 - \pi)G(F(1, \sigma, \xi); \psi) & y = 0 \\ (1 - \pi)[G(F(y + 1, \sigma, \xi); \psi) - G(F(y, \sigma, \xi); \psi)] & y \in \mathbb{N} \end{cases} \quad (3.1)$$

where $0 \leq \pi \leq 1$ is the mixing proportion, determining from which state Y is generated. In the following, we coin (3.1) as the ZIDEGPD model.

Suppose now that $\mathbf{x} \in \mathbb{R}^q$ is a vector of covariates measured with Y . In a continuous framework, the inclusion of flexible forms of dependence of extreme values on covariates is well established (Davison and Smith, 1990). Chavez-Demoulin and Davison (2005) used the Generalized Additive Model (GAM) (Hastie and Tibshirani, 1990) to flexibly estimate GPD parameters. More recently, by coupling GAM forms with penalized likelihood, Youngman (2019) modeled threshold exceedances with GPD parameters of GAM forms.

By allowing the parameters to depend on covariates, we extend the pmf (3.1) to the zero-inflated count regression setting. More specifically, we identify the vector of parameters (ξ, σ, ψ, π) with $\theta = (\theta_1, \dots, \theta_d)$. The parameters of the distribution of Y can depend on the covariates \mathbf{x} , i.e. $\theta(\mathbf{x}) = (\theta_1(\mathbf{x}), \dots, \theta_d(\mathbf{x}))$. To relate the distribution parameters $(\theta_1(\mathbf{x}), \dots, \theta_d(\mathbf{x}))$ to the covariates, we consider additive

predictors of the form

$$\eta_i(\mathbf{x}) = s_{i1}(\mathbf{x}) + \cdots + s_{iJ_i}(\mathbf{x}), \quad i = 1, \dots, d, \quad (3.2)$$

where $s_{i1}(\cdot), \dots, s_{iJ_i}(\cdot)$ are smooth functions of the covariates \mathbf{x} . The predictors are linked to the parameters via known monotonic and twice differentiable link functions $h_i(\cdot)$.

$$\theta_i(\mathbf{x}) = h_i(\eta_i(\mathbf{x})), \quad i = 1, \dots, d. \quad (3.3)$$

For instance, we use the following linking functions for the model (i) associated with $G(u; \psi) = u^\kappa$. The parameters can be written as

$$\xi(\mathbf{x}) = \exp(\eta_\xi(\mathbf{x})), \quad \sigma(\mathbf{x}) = \exp(\eta_\sigma(\mathbf{x})), \quad \kappa(\mathbf{x}) = \exp(\eta_\kappa(\mathbf{x})), \quad \pi(\mathbf{x}) = \exp\left(\frac{\eta_\pi(\mathbf{x})}{1 + \eta_\pi(\mathbf{x})}\right).$$

The functions $s_{ij}(\cdot)$ in (3.2) are approximated by a set of K_{ij} basis functions $\{B_{k,ij}(\mathbf{x}) \mid k = 1, \dots, K_{ij}\}$, namely

$$s_{ij}(\mathbf{x}) = \sum_{k=1}^{K_{ij}} \beta_{ij,k} B_k(\mathbf{x}). \quad (3.4)$$

The basis functions can be of different types (see Wood, 2017, for instance). The basis function expansions can be written as $s_{ij}(\mathbf{x}) = \mathbf{t}_{ij}(\mathbf{x})^T \boldsymbol{\beta}_{ij}$ where $\mathbf{t}_{ij}(\mathbf{x})$ is still a vector of transformed covariates that depends on the basis functions and $\boldsymbol{\beta}_{ij} = (\beta_{ij,1}, \dots, \beta_{ij,K_{ij}})^T$ is a parameter vector to be estimated.

The penalized maximum likelihood estimation (MLE) method is used to estimate the parameters of the proposed models. More precisely, let y_1, \dots, y_n be n independent observations from (2.6) and $\mathbf{x}_1, \dots, \mathbf{x}_n$ the related covariates. The log-likelihood function is given by

$$\begin{aligned} l(\boldsymbol{\beta}) &= \sum_{i=1}^n I_{\{0\}}(y_i) \log [\pi(\mathbf{x}_i) + (1 - \pi(\mathbf{x}_i))G(F(1; \sigma(\mathbf{x}_i), \xi(\mathbf{x}_i)); \psi(\mathbf{x}_i))] \\ &\quad + \sum_{i=1}^n (1 - I_{\{0\}}(y_i)) \log(1 - \pi(\mathbf{x}_i)) \times \\ &\quad [G(F(y_i + 1; \sigma(\mathbf{x}_i), \xi(\mathbf{x}_i)); \psi(\mathbf{x}_i)) - G(F(y_i; \sigma(\mathbf{x}_i), \xi(\mathbf{x}_i)); \psi(\mathbf{x}_i))], \end{aligned} \quad (3.5)$$

where $I_A(\cdot)$ is the indicator function of the set A . To ensure regularization of the functions $s_{ij}(\mathbf{x})$ so-called penalty terms are added to the objective log-likelihood function. Usually, the penalty for each function $s_{ij}(\mathbf{x})$ is a quadratic penalty $\boldsymbol{\beta}_{ij}^T \mathbf{P}_{ij}(\boldsymbol{\lambda}_{ij}) \boldsymbol{\beta}_{ij}$ where $\mathbf{P}_{ij}(\boldsymbol{\lambda}_{ij})$ is a known semi-definite matrix and the vector $\boldsymbol{\lambda}_{ij}$ regulates the amount of smoothing needed for the fit. A special case that we use in the real data application is when $\mathbf{P}_{ij}(\boldsymbol{\lambda}_{ij}) = \lambda_{ij} \mathbf{P}_{ij}$, for a scalar $\lambda_{ij} > 0$ and a semi-definite matrix \mathbf{P}_{ij} . The entries of the penalty matrix \mathbf{P}_{ij} are the integrals of the products of the second derivatives of pairs of cubic spline functions, see Wood (2011, Section 5.3)

for more details. The penalized log-likelihood function for the latter models reads:

$$l_p(\boldsymbol{\beta}) = l(\boldsymbol{\beta}) - \frac{1}{2} \sum_{i=1}^d \sum_{j=1}^{J_i} \lambda_{ij} \boldsymbol{\beta}_{ij}^T \mathbf{P}_{ij} \boldsymbol{\beta}_{ij}. \quad (3.6)$$

We apply the restricted maximum likelihood (REML) approach to estimate $\boldsymbol{\beta}_{ij}$ and λ_{ij} following Wood (2011). Our current implementation exploits the R package `evgam` (Youngman, 2022) and adds two new families of distributions named `degpd` and `zidegpd`. Note that within the `degpd` family, all four models for $G(\cdot; \psi)$ have been implemented. By contrast, for the `zidegpd` family, only models (i), (ii), and (iii) have been implemented, since the `evgam` package allows the simultaneous estimation of only five distribution parameters $\boldsymbol{\theta}(\mathbf{x}) = (\theta_1(\mathbf{x}), \dots, \theta_5(\mathbf{x}))^T$.

4 Avalanche data example

Let us return to the example introduced with Figure 1. The data comes from the *Enquête Permanente sur les Avalanches*. This survey collects avalanche data from the French Alps and has monitored about 3900 routes by local observers since the beginning of the 20th century (see Mougin, 1922; Evin et al., 2021). Quantitative (run-out elevations, deposit volumes, etc.) and qualitative (flow regime, snow quality, etc.) information is collected for each event.

Report quality of avalanches observed by local observers varies over time and space. For example, an avalanche event may be recorded a few days after it occurs, and the estimated day of the event by the observer may be approximate. It is too restrictive to select only the events for which the day is known, as too many events would be lost. However, if the avalanche occurred several days before the observation, including the observation could lead to biased analysis, such as an inaccurate count of avalanches for that day and difficulty in linking the event to snow and meteorological conditions (Evin et al., 2021). Therefore, we only consider avalanche events that occurred within three days before of their observation.

Natural avalanche activity is also uncertain because records tend to record paths visible from valleys, so the high-elevation activity may be underestimated. Avalanches are usually caused by severe storms that bring high snowfalls coupled with snow drifting, but substantial variations of environmental covariates causing snow melt and/or fluctuations of the freezing point can also be involved. For instance, precipitation amounts, air temperatures during storms, and prior snow stratigraphy influence avalanche types and frequency. Although overall avalanche frequency is likely to decrease globally (Strapazzon et al., 2021), more extreme environmental conditions during winter storms can cause more intensive avalanche events. For instance, a shallow snow-pack and warmer temperatures have become increasingly influential

in the Alps. Since extreme events have potentially terrible consequences, it is crucial to anticipate future avalanche activity in the short-term and long-term management, possibly relating this activity to environmental variables.

In our example the total number of avalanches in the Haute Maurienne massif was considered within a three-day observation window between January 1982 and April 2021. Only days between October 15th and May 15th were included. A total of 2839 observations were made during this period [Dkengne et al. \(2016\)](#).

The related environmental variables are maximum wind speed (WS) at 10 meters in m/s, relative humidity (RH) at 2 meters in percentage, precipitation (PREC) in mm per day, maximum (MxT) and minimum (MnT) temperatures at 2 meters in °C. Environmental covariates have been downloaded from <https://power.larc.nasa.gov/data-access-viewer/> by specifying latitude and longitude information. Three-day summaries are calculated by averaging RH and PREC and taking the maximum for MxT and WS and the minimum for MnT.

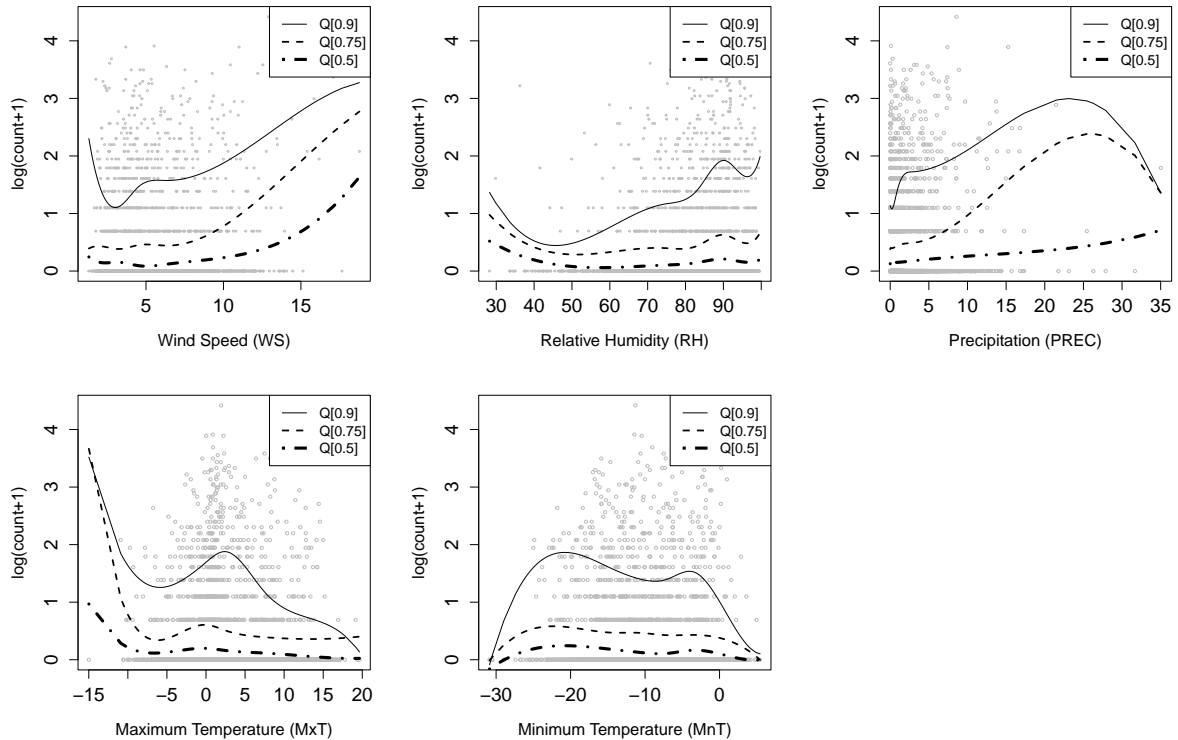


Figure 3: Scatterplots of avalanche counts (log scale) versus environmental variables. The smoothing lines are obtained by jittering the count and fitting three non-parametric regression quantile models at the 0.5, 0.75, and 0.90 quantiles.

The scatterplots of avalanche counts (log scale) versus environmental variables (Figure 3) do not highlight clear relationships, mainly masked by the presence of many zeros.

	WS	RH	PREC	MxT	MnT
WS	1.00	0.17	0.27	-0.07	-0.18
RH	0.17	1.00	0.43	-0.39	-0.24
PREC	0.27	0.43	1.00	0.01	0.11
MxT	-0.07	-0.39	0.01	1.00	0.86
MnT	-0.18	-0.24	0.11	0.86	1.00

Figure 4: Correlation between environmental variables.

Figure 4 shows the correlation plot between the covariates, highlighting that maximum temperature (MxT) and minimum temperature (MnT) are positively strongly correlated. At the same time, precipitation (PREC) has no significant correlation with MxT. On the other hand, relative humidity (RH) has a low and moderate positive correlation with wind speed (WS) and PREC, while it has a weak negative correlation with temperature variables. Furthermore, wind speed and precipitation have a weak correlation with minimum and maximum temperature variables.

We fit the ZIDEGPD model under specifications (i), (ii), and (iii) for $G(u; \psi)$, and a backward variable selection procedure based on AIC is performed for selecting the covariates. The shape parameter ξ is kept constant to avoid numerical instability in the estimation of ψ . The results of the selection are given in Table 1, where $s(\cdot)$ indicates the smoothed predictor.

To assess the overall adequacy of ZIDEGPD model, we used the randomized residuals (Dunn and Smyth, 1996) defined as

$$r_i = \Phi^{-1} \left((1 - u_i)F(y_i - 1; \hat{\boldsymbol{\theta}}(\mathbf{x}_i)) + u_i F(y_i; \hat{\boldsymbol{\theta}}(\mathbf{x}_i)) \right)$$

where Φ^{-1} is the inverse of CDF of a standard Gaussian distribution function, u_i is drawn from a uniform distribution, and $F(\cdot; \boldsymbol{\theta})$ is the CDF of the current model. Aside

Table 1: Selected predictors for ZIDEGPD model ^a

Model	$\log \sigma(\mathbf{x})$	$\log \kappa(\mathbf{x})$	$\log \delta(\mathbf{x})$	$\text{logit}(\pi(\mathbf{x}))$
(i)	$\mathbf{s}(\text{WS}) + \mathbf{s}(\text{MxT}) + \mathbf{s}(\text{PREC})$	$\mathbf{s}(\text{RH})$	-	$\mathbf{s}(\text{MxT})$
(ii)	$\mathbf{s}(\text{WS}) + \mathbf{s}(\text{MxT}) + \mathbf{s}(\text{RH})$	-	cst	$\mathbf{s}(\text{MxT}) + \mathbf{s}(\text{PREC})$
(iii)	$\mathbf{s}(\text{WS}) + \mathbf{s}(\text{MxT}) + \mathbf{s}(\text{RH}) + \mathbf{s}(\text{PREC})$	cst	cst	$\mathbf{s}(\text{MxT})$

^a cst denotes the constant parameters.

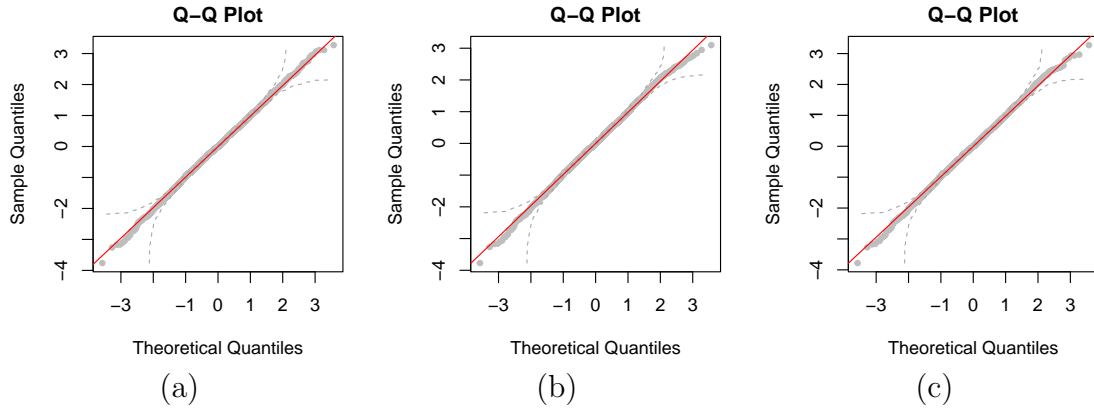


Figure 5: Q-Q plot of randomized residuals of different ZIDEGPD models: (a) Model (i); (b) Model (ii); (c) Model (iii). The dotted lines show the 95% point-wise confidence intervals.

from sampling variation in the parameter estimates, randomized residuals should follow a standard Gaussian distribution if the model is correctly identified. The randomized residuals derived from our proposed model show no apparent departure from normality (Figure 5) and fit both tails (lower and upper) correctly when compared to the existing models in the literature, as shown in Figure 1.

In Table 2, we report the final values of the AIC and values of the Anderson-Darling test for the normality of the randomized residuals. Since looking at just one realization can be misleading, we compute the Anderson-Darling test statistic for 1,000 realizations of the randomized residuals, take the median of the test statistic, and evaluate the corresponding P-value. As we can see, Model (i) is slightly preferred in terms of goodness of fit, and Table 3 summarizes the output of the fitting procedure for this model. Note that the smoothing terms are all statistically significant, in particular, there is no indication for $\kappa(x) = 1$, i.e. when the DEGP reduces to a DGP in the mixture.

A broad interpretation of the results (see also Figure 6) is that meteorological conditions have an increasing effect on avalanche occurrence, except for the maximum of temperature, which shows a non-monotonic pattern around 0 degrees. Only the temperature affects the mixing proportion π , with an increment for temperatures above

Table 2: AIC values and Anderson-Darling normality test statistics of the different ZIDEGPD models.

Model	AIC	Statistics	P-value
(i)	6121	0.287	0.623
(ii)	6151	0.367	0.432
(iii)	6152	0.315	0.543

Table 3: Estimated coefficients and smooth terms for ZIDEGPD model (i) fitted to avalanches data. The p-values of the smoothed terms $\mathbf{s}(\cdot)$ indicate the significance of their presence.

ZIDEGPD with $G(u; \psi) = u^\kappa$				
Constant terms				
	Estimate	Std Error	t value	P-value
$\log(\kappa)$	-1.62	0.22	-7.42	<0.001
$\log(\sigma)$	1.62	0.15	10.71	<0.001
$\log(\xi)$	-1.63	0.43	-3.83	<0.001
$\text{logit}(\pi)$	-1.25	0.73	-1.71	0.044
Smooth terms for $\log(\kappa)$				
	edf	max.df	Chi.sq	$\Pr(> t)$
$\mathbf{s}(\text{RH})$	1.02	9	11.78	<0.001
Smooth terms for $\log(\sigma)$				
	edf	max.df	Chi.sq	$\Pr(> t)$
$\mathbf{s}(\text{WS})$	1.88	7	14.80	0.005
$\mathbf{s}(\text{MxT})$	4.81	7	26.76	<0.001
$\mathbf{s}(\text{PREC})$	1.10	4	7.16	0.008
Smooth terms for $\text{logit}(\pi)$				
	edf	max.df	Chi.sq	$\Pr(> t)$
$\mathbf{s}(\text{MxT})$	1.16	7	7.32	0.013

0°C. The value of κ increases with relative humidity, which is associated (see Figure 2-(a)) with an increasing probability of extreme avalanches. Our results are in agreement with those of the existing literature (Dreier et al., 2013) where, for example, snow surface, air temperatures, and changes in snow height and relative humidity strongly influenced snow slides in spring periods.




Figure 6: Estimated non-parametric effects of covariates in parameters of ZIDEGPD (i).

5 Conclusions

There are many examples of distributions for count data. Motivated by the avalanche real data example, we attempt to properly model the lower limit of the distribution and the upper tail without neglecting the bulk part of the data.

The family of distributions deriving from an efficient choice of the function $G(u; \psi)$ in the DEGP offers great flexibility in their behavior in both tails. Additionally, we apply mixture models within this framework to further enhance modeling capabilities, particularly in addressing excess zeros commonly observed in such datasets.

We have developed and implemented R functions to fit DEGP and ZIDEGPD pa-

rameters in GAM forms that allow for non-identically distributed discrete extremes. These functions use the functions in `evgam` (Youngman, 2022). As a result, it is possible to consider non-additive model formulations and fit them using thin-plate splines, which are particularly attractive for modeling isotropic spatial dependence, or tensor products of splines for modeling interactions between covariates.

In terms of application, the variability of avalanche activity has been statistically related to other environmental variables (e.g., temperature, wind, precipitation, and humidity). Although this model does not allow for short-term predictions, it does allow for the association of specific weather conditions with avalanche risk by allowing experts to account for possible nonlinearity. Indeed the GAMs proposed in this study allow parametric and non-parametric functional forms, which would most likely be required for larger data sets. Compared to other competing models available in the literature our proposed models are more flexible in estimating both tail behavior and achieving a better fit for avalanche data when environmental conditions are considered as covariates.

Acknowledgements

The authors thank Benjamin Youngman for helpful discussions on the R code and Nicolas Eckert for providing and explaining the avalanche data. Touqeer Ahmad acknowledges support from the Région Bretagne through project SAD-2021- MaEVa. Philippe Naveau acknowledges the support of the French Agence Nationale de la Recherche (ANR) under reference ANR-Melody (ANR-19-CE46-0011). Part of this work was also supported by 80 PRIME CNRS-INSU, ANR-20-CE40-0025-01 (T-REX project), and the European H2020 XAIDA (Grant agreement ID: 101003469).

Code availability

Full source code for the proposed models with a simple running example is available at <https://github.com/touqeerahmadunipd/degpd-and-zidegpd>.

References

- Balkema, A. A. and De Haan, L. (1974). Residual life time at great age. *The Annals of Probability*, **2**, 792–804.
- Beirlant, J., Goegebeur, Y., Segers, J., and Teugels, J. L. (2004). *Statistics of Extremes: Theory and Applications*. John Wiley & Sons, New York.

- Buddana, A. and Kozubowski, T. J. (2014). Discrete Pareto distributions. *Economic Quality Control*, **29**, 143–156.
- Carreau, J. and Bengio, Y. (2009). A hybrid Pareto model for asymmetric fat-tailed data: the univariate case. *Extremes*, **12**, 53–76.
- Chavez-Demoulin, V. and Davison, A. C. (2005). Generalized additive modelling of sample extremes. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **54**, 207–222.
- Coles, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer, New York.
- Daouia, A., Stupfler, G., and Usseglio-Carleve, A. (2023). Extreme value modelling of SARS-CoV-2 community transmission using discrete generalized Pareto distributions. *Royal Society Open Science*, **10**.
- Davison, A. C. and Smith, R. L. (1990). Models for exceedances over high thresholds. *Journal of the Royal Statistical Society: Series B (Methodological)*, **52**, 393–425.
- de Carvalho, M., Pereira, S., Pereira, P., and de Zea Bermudez, P. (2022). An extreme value bayesian lasso for the conditional left and right tails. *Journal of Agricultural, Biological and Environmental Statistics*, **27**, 222–239.
- Dkengne, P. S., Eckert, N., and 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.
- Dreier, L., Mitterer, C., Feick, S., Harvey, S., et al. (2013). The influence of weather on glide-snow avalanches. In *Proceedings, International Snow Science Workshop, France, Grenoble*, volume 2951.
- Dunn, P. K. and Smyth, G. K. (1996). Randomized quantile residuals. *Journal of Computational and Graphical Statistics*, **5**, 236–244.
- Eastoe, E. F. and Tawn, J. A. (2009). Modelling non-stationary extremes with application to surface level ozone. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **58**, 25–45.
- Evin, G., Sielenou, P. D., Eckert, N., Naveau, P., Hagenmuller, P., and Morin, S. (2021). Extreme avalanche cycles: return levels and probability distributions depending on snow and meteorological conditions. *Weather and Climate Extremes*, **33**, 100344.
- Fisher, R. A. and Tippett, L. H. C. (1928). Limiting forms of the frequency distribution of the largest or smallest member of a sample. *Mathematical Proceedings of the Cambridge Philosophical Society*, **24**, 180–190.

- Frigessi, A., Haug, O., and Rue, H. (2002). A dynamic mixture model for unsupervised tail estimation without threshold selection. *Extremes*, **5**, 219–235.
- Hastie, T. and Tibshirani, R. (1990). *Generalized Additive Models*. Chapman & Hall, CRC Press, New York.
- Hitz, A. S., Davis, R. A., and Samorodnitsky, G. (2024). Discrete extremes. *Journal of Data Science*, pages 1–13.
- Kneib, T., Silbersdorff, A., and Säfken, B. (2023). Rage against the mean – a review of distributional regression approaches. *Econometrics and Statistics*, **26**, 99–123.
- Kozubowski, T. J., Panorska, A. K., and Forister, M. L. (2015). A discrete truncated Pareto distribution. *Statistical Methodology*, **26**, 135–150.
- Krishna, H. and Pundir, P. S. (2009). Discrete Burr and discrete Pareto distributions. *Statistical Methodology*, **6**, 177–188.
- Lambert, D. (1992). Zero-inflated poisson regression, with an application to defects in manufacturing. *Technometrics*, **34**, 1–14.
- MacDonald, A., Scarrott, C. J., Lee, D., Darlow, B., Reale, M., and Russell, G. (2011). A flexible extreme value mixture model. *Computational Statistics & Data Analysis*, **55**, 2137–2157.
- Mougin, P. (1922). Avalanches in Savoy, vol. iv. *Ministry of Agriculture, General Directorate of Water and Forests, Department of Great Hydraulic Forces, Paris*.
- Naveau, P., Huser, R., Ribereau, P., and Hannart, A. (2016). Modeling jointly low, moderate, and heavy rainfall intensities without a threshold selection. *Water Resources Research*, **52**, 2753–2769.
- Papastathopoulos, I. and Tawn, J. A. (2013). Extended generalised Pareto models for tail estimation. *Journal of Statistical Planning and Inference*, **143**, 131–143.
- Pickands, J. (1975). Statistical inference using extreme order statistics. *The Annals of Statistics*, **3**, 119–131.
- Ranjbar, S., Cantoni, E., Chavez-Demoulin, V., Marra, G., Radice, R., and Jatou, K. (2022). Modelling the extremes of seasonal viruses and hospital congestion: The example of flu in a swiss hospital. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **71**, 884–905.
- Rigby, R. A. and Stasinopoulos, D. M. (2005). Generalized additive models for location, scale and shape. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **54**, 507–554.
- Stasinopoulos, M. D., Rigby, R. A., and Bastiani, F. D. (2018). GAMLSS: A distributional regression approach. *Statistical Modelling*, **18**, 248–273.

- Stein, M. L. (2021). A parametric model for distributions with flexible behavior in both tails. *Environmetrics*, **32**, e2658.
- Strapazzon, G., Schweizer, J., Chiambretti, I., Brodmann Maeder, M., Brugger, H., and Zafren, K. (2021). Effects of climate change on avalanche accidents and survival. *Frontiers in Physiology*, **12**, 639433.
- Tencaliec, P., Favre, A.-C., Naveau, P., Prieur, C., and Nicolet, G. (2020). Flexible semiparametric generalized Pareto modeling of the entire range of rainfall amount. *Environmetrics*, **31**, e2582.
- Wedderburn, R. W. M. (1974). Quasi-likelihood functions, generalized linear models, and the Gauss-Newton method. *Biometrika*, **61**, 439–447.
- Wood, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **73**, 3–36.
- Wood, S. N. (2017). *Generalized Additive Models: an Introduction with R*. Chapman and Hall/CRC, New York, 2nd edition.
- Youngman, B. D. (2019). Generalized additive models for exceedances of high thresholds with an application to return level estimation for US wind gusts. *Journal of the American Statistical Association*, **114**, 1865–1879.
- Youngman, B. D. (2022). evgam: An R package for Generalized Additive Extreme Value Models. *Journal of Statistical Software*, **103**, 1–26.