

From the power law to extreme value mixture distributions

Clement Lee

School of Mathematics, Statistics and Physics, Newcastle University
and

Emma F. Eastoe and Aiden Farrell

Department of Mathematics and Statistics, Lancaster University

January 18, 2024

Abstract

The power law is useful in describing count phenomena such as network degrees and word frequencies. With a single parameter, it captures the main feature that the frequencies are linear on the log-log scale. Nevertheless, there have been criticisms of the power law, for example that a threshold needs to be pre-selected without its uncertainty quantified, that the power law is simply inadequate, and that subsequent hypothesis tests are required to determine whether the data could have come from the power law. We propose a modelling framework that combines two different generalisations of the power law, namely the generalised Pareto distribution and the Zipf-polylog distribution, to resolve these issues. The proposed mixture distributions are shown to fit the data well and quantify the threshold uncertainty in a natural way. A model selection step embedded in the Bayesian inference algorithm further answers the question whether the power law is adequate.

Keywords: degree distribution; generalised Pareto; polylogarithm; threshold uncertainty; Markov chain Monte Carlo; Bayesian model selection

1 Introduction

In statistics, the power law is a principle by which the relative change in the size of a random variable is proportional to the relative change in its frequency. The continuous distribution associated with the power law is the Pareto distribution, of which the probability density function (PDF) is

$$f(z) = \frac{(\alpha - 1)}{u} \left(\frac{z}{u}\right)^{-\alpha}, \quad z > u, \quad (1)$$

where $u > 0$ is the threshold, and $\alpha > 1$ is the exponent parameter. The analogous distribution for discrete data is the Zipf distribution, with probability mass function (PMF)

$$p(x) = \frac{x^{-\alpha}}{\zeta(\alpha, u + 1)}, \quad x = u + 1, u + 2, \dots, \quad (2)$$

where u is a non-negative integer, and $\zeta(\alpha, z) = \sum_{i=0}^{\infty} (z + i)^{-\alpha}$ is the Hurwitz zeta function. Whilst the form of this distribution with $u = 0$ has many other names including the zeta distribution and the discrete Pareto distribution (Johnson et al., 2005), we use the name Zipf distribution for consistency with Valero et al. (2022). The Zipf distribution has been applied to data in various fields including quantitative linguistics (Bell et al., 2012), casualty numbers (Bohorquez et al., 2009, Friedman, 2015, Gillespie, 2017), and city sizes (Newman, 2005, Clauset et al., 2009).

In this paper, we focus on a specific class of discrete data which is routinely associated with the power law – the (in-)degrees of a network. Examples include networks of links on the World-Wide Web (Barabási and Albert, 1999, Albert et al., 1999, Faloutsos et al., 1999), social networks (Lee and Oh, 2014, Varga, 2015), coauthorship and citation networks (Price, 1976, Thelwall and Wilson, 2014, Arroyo-Machado et al., 2020), and retweet counts (Bhamidi et al., 2015, Mathews et al., 2017). Networks for which the degree distribution follows a power law are described as scale-free.

When studying the degree distribution of a network, the behaviour of the large (in-)degrees is of most interest because this can identify the generative mechanism of the network. Specifically, the popular preferential attachment (PA) generative model (Barabási and Albert, 1999), which extends the cumulative advantage process (Price, 1976), generates networks for which the right tail of the degree distribution follows the power law. Therefore, determining whether the empirical right tail of the degree distribution of a network follows a power law can inform the choice of model for the network generating mechanism.

Despite the popularity of the PA model, many have argued that application of the power law to degree distributions is inappropriate, that scale-free networks are limiting and that the PA model is inadequate. For example, Stumpf et al. (2005) found that sampling from scale-free networks does not result in (smaller) scale-free networks. Stumpf and Porter (2012) questioned the ubiquity of the power law in different disciplines, whilst Broido and Clauset (2019) found that scale-free networks are rare in reality, and for most networks the log-normal distribution is sufficient. While Voitalov et al. (2019) criticised the latter, the findings have been supported by, for example, Thelwall and Wilson (2014), Sheridan and Onodera (2018), and Arroyo-Machado et al. (2020).

This paper aims to address these, and other, concerns about the use of the power law to model degree distributions. To motivate our work, we examine some real-world examples of network degree distributions using two popular graphical diagnostics. From Equation 2 the logarithm of the Zipf distribution, $\log(p(x))$, is linear with $\log(x)$. Hence the power law is a suitable model for a degree distribution if the empirical frequency, or equivalently the empirical survival function, has a linear relationship with the data on the log-log scale. Figure 1 displays both diagnostics for four examples, two from the KONECT database (Kunegis, 2013) and two from Valero et al. (2022). First observe that, for all data sets,

a large proportion of the nodes in the network have very small degrees. For example, for both data sets on the left, the values $x = 1, \dots, 10$ make up 81% of the data. In other words, if the right tail of the degree distribution were defined to be all values above a high threshold, then it would consist of almost all the unique values in the data set.

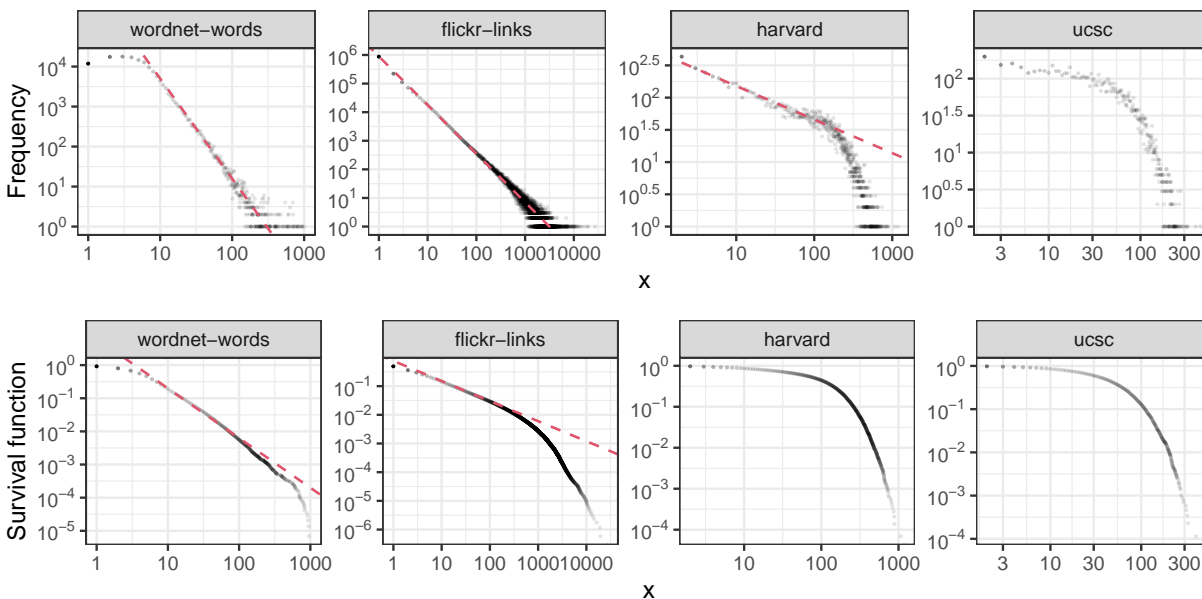


Figure 1: Empirical frequencies (top) and survival function (bottom) on the log-log scale of four data sets. Barring the rightmost data set, a red dashed line with slope $-\alpha$ in the top plot implies a dashed line with slope $-(\alpha - 1)$ in the bottom plot, and suggests a power law with exponent α .

Now consider the linearity of the plots. The “wordnet-words” data shows piecewise linearity, with the left and right tails showing different behaviour. In this case, it might be appropriate to fit the Zipf distribution to the largest degrees, however this requires a subjective choice of a cut-off point, or threshold, with which to identify the largest degrees which, in turn, will impact inference. Further, by construction, the models are fitted using different subsets of the data meaning standard model selection tests cannot be used to identify the threshold.. Clauset et al. (2009) and Gillespie (2015) suggested using the

Kolmogorov-Smirnov statistic, but this requires additional procedures beyond fitting the Zipf distribution whilst not resolving the loss of information.

The “flickr-links” data is a case in which the Zipf distribution might fit the body of the data well but would be insufficiently flexible for the right tail. This data also clearly demonstrates the utility of the survival plot over the frequency plot. That the right tail should be separately modelled aligns with the work of Voitalov et al. (2019), who suggested the use of regularly varying distributions in such cases. Since such distributions have heavy right tails, this approach struggles with data with light tails and, further, still requires a subset of the data to be identified prior to performing inference.

The body and tail of the “harvard” data also differ. In this case, when the right tail is removed, a truncated Zipf distribution with $\alpha \leq 1$ is appropriate. Aside from the lack of fit to the right tail, such a power law is not appropriate for the whole support of the data since the resulting distribution does not have a proper PMF. Lastly, as shown by the “ucsc” data, the power law may simply not be appropriate. In this case, the survival function is neither fully, nor piecewise, linear.

In summary, the Zipf distribution is too simplistic to satisfactorily capture the empirical degree distribution of many real-world data sets, and even when the power law is appropriate for a subset of the data, there is neither a systematic, objective way of choosing this subset nor a way to quantify the uncertainty of this choice. We address both issues by proposing a modelling and inference framework that uses a mixture of discrete distributions to accommodate survival curves with a piecewise nature. While Jung and Phoa (2021) proposed a similar approach, their model is a mixture of Zipf distributions with the additional restriction that the same exponent parameter α be used for all components, resulting in very limited flexibility. In contrast, we use a mixture of Zipf-type and extreme

value distributions, which have disjoint supports and model the body and the tail(s) separately. More importantly, analogous to existing continuous counterparts in the literature, our approach allows estimation of the threshold(s) between the component distributions. Using a Bayesian inference approach permits both estimation of threshold uncertainty and model selection. The latter embedded naturally in the inference algorithm, thus removing the need for further hypothesis testing.

The rest of the paper is as follows. In Section 2, distributions related to the power law and extreme value theory are introduced, along with the characterisation of the asymptotic tail behaviour of the former. Two of these distributions form the mixture model, which is proposed in Section 3. The inference algorithm is outlined in Section 4. Applications to real-world data are presented in Section 5. The discussion in Section 6 concludes the paper.

2 Distributions and properties

2.1 The truncated Zipf-polylog distribution

The *truncated* Zipf-polylog(α, θ, u, w) (TZP) distribution has the PMF

$$p_{\text{TZP}}(x; \alpha, \theta, u, w) = \frac{x^{-\alpha}\theta^x}{\sum_{k=w+1}^u k^{-\alpha}\theta^k}, \quad x = w + 1, w + 2, \dots, u, \quad (3)$$

where $\theta \in (0, 1]$, u and w are integers such that $u > w \geq 0$, and $\alpha \in \mathbb{R}$ when $\theta \in (0, 1)$ or $u < \infty$, and $\alpha > 1$ when $\theta = 1$ and $u = \infty$. Several special cases of this distribution are summarised in Figure 2, including the Zipf-polylog (ZP) distribution (Valero et al., 2022). In what follows, the *polylog* distribution refers to the ZP distribution when $\theta \in (0, 1)$, and consequently the $\text{ZP}(\alpha, \theta, w)$ distribution is the disjoint union of the $\text{Zipf}(\alpha, w)$ and

2.2 From the GP distribution to its discrete version

We next look at how the GP distribution generalises the Pareto distribution, how it is used as a statistical model, and its discretisation. We begin with its PDF

$$g(z) = \begin{cases} \frac{1}{\sigma_0} \left[1 + \xi \frac{z - \mu}{\sigma_0} \right]_+^{-1/\xi - 1}, & \xi \neq 0, \\ \frac{1}{\sigma_0} \exp \left[-\frac{z - \mu}{\sigma_0} \right], & \xi = 0, \end{cases}$$

where $A_+ := \max\{A, 0\}$, μ is the location parameter, $\sigma_0 > 0$ is the scale parameter, and ξ is the tail index or shape parameter. The support is $x \geq \mu$ when $\xi \geq 0$, and $\mu \leq x \leq \mu - \sigma_0/\xi$ when $\xi < 0$. There is no loss in generality in working primarily with the upper expression as the lower expression is its limit as $\xi \rightarrow 0$. If a random variable Z follows the $\text{GP}(\mu, \sigma_0, \xi)$ distribution exactly, the distribution of $Z|Z > u$ has the PDF

$$g_u(z) = \frac{1}{\sigma_0 + \xi(u - \mu)} \left[1 + \frac{\xi(z - u)}{\sigma_0 + \xi(u - \mu)} \right]_+^{-1/\xi - 1}, \quad (4)$$

which is the $\text{GP}(u, \sigma_0 + \xi(u - \mu), \xi)$ distribution. To see how the GP distribution generalises the Pareto distribution, note that when $\xi > 0$ and $\sigma_0 = \xi\mu$, $g_u(z)$ is equivalent to the density in Equation 1 with $\alpha = 1/\xi + 1$. In practice, for identifiability, the GP distribution in Equation 4 is re-parametrised as (ϕ_u, σ, ξ) where

$$\phi_u := \Pr(Z > u) = \left[1 + \xi \frac{u - \mu}{\sigma_0} \right]_+^{-1/\xi}$$

is the exceedance probability, and $\sigma = \sigma_0 - \xi\mu$. For this and other practical considerations when fitting the GP distribution to real-world data, see, for example, Coles (2001).

The practical importance of the GP distribution comes from the work of Pickands (1975), who showed that for $Z \sim F$ where F is an arbitrary continuous CDF and high u , the GP distribution approximates the conditional tail $\Pr(Z > z|Z > u)$. Further, the GP distribution is the only distribution with this capability. The approximation follows

from asymptotic results which show that for almost all absolutely continuous distributions F , the conditional distribution in Equation 4 holds exactly in the limit as $u \rightarrow \infty$. These results support the use of the GP distribution in data analysis as a model for observations above a high threshold. We will look at how this threshold is chosen in Section 3.

As the GP distribution is a continuous distribution, it is not appropriate to apply it directly to discrete data. Whilst Voitalov et al. (2019) proposed adding uniform noise to integer-valued data before estimating ξ using the continuous distribution, a more elegant solution is to use the integer-GP (IGP) distribution (Prieto et al., 2014, Rohrbeck et al., 2018), the discretised version of the GP distribution. Specifically, if Z follows the GP distribution with PDF $g(z)$, then $X = \lceil Z \rceil$ follows the IGP distribution with the same parameters and has PMF $p(x) = \int_{x-1}^x g(z)dz$. The case where $X = \lfloor Z \rfloor$ has been considered by Hitz et al. (2017). The relationships between the Pareto, GP and IGP distributions are summarised in the rightmost column of Figure 2. Note that the IGP distribution has no direct relationship with the TZP distribution, even though both are discrete extensions of power law distributions.

2.3 Domain of attraction of Zipf-polylog

The behaviour of the right tail of a continuous probability distribution F_0 can be characterised by the maximum domain of attraction (DoA) of the distribution. Specifically, F_0 is in the DoA of an extreme value distribution H if there exists $a_n > 0$, $b_n \in \mathbb{R}$ such that $\lim_{n \rightarrow \infty} |F_0^n(a_n x + b_n) - H(x)| = 0$, where H must be a negative Weibull, Gumbel, or Fréchet distribution. There is a one-to-one connection between this result and the GP distribution tail approximation: when the right tail of F_0 is approximated by the GP distribution with $\xi < 0$, $\xi = 0$, and $\xi > 0$, then F_0 is in the DoA of the negative Weibull,

Gumbel, and Fréchet distributions, respectively.

While the unified DoA of the three extreme value distributions covers a wide range of common continuous distributions, many common discrete distributions, including the geometric, do not belong to a DoA according to the definition above. To address this, Shimura (2012) introduced the concept of recovery to a DoA. If a discrete distribution F is the discretisation of a continuous distribution F_0 and F_0 is in a DoA, then F is said to be recoverable to the same DoA. For example, the geometric distribution is a discretisation of the exponential distribution, which is in the Gumbel DoA, therefore the geometric distribution is recoverable to the Gumbel DoA.

As the ZP distribution is the disjoint union of the polylog and Zipf distributions, understanding the right tail behaviour of the ZP distribution requires us to first characterise the tail behaviour of both these distributions. Specifically, the DoA to which each of the polylog and Zipf distributions is recoverable can be used to determine the behaviour of the equivalent IGP tail approximation.

For the polylog distribution, the key quantity for recovery is

$$\Omega(F, x) := \left(\log \frac{S(x+1)}{S(x+2)} \right)^{-1} - \left(\log \frac{S(x)}{S(x+1)} \right)^{-1},$$

where $S(x) = 1 - F(x)$ is the survival function of the discrete distribution concerned. Shimura (2012) showed that if $\lim_{x \rightarrow \infty} \Omega(F, x) = 0$, then F is (uniquely) recoverable to the Gumbel DoA. This turns out to be the case for the polylog(α, θ, w) distribution; see Appendix A.1 for the proof. Consequently, the IGP distribution that approximates the polylog(α, θ, w) distribution has $\xi = 0$ always.

The Zipf distribution requires the result that if $\lim_{x \rightarrow \infty} \Omega(F, x) = \xi > 0$, then the discrete distribution is in the DoA of the Fréchet distribution with positive tail index ξ (Shimura, 2012). For the Zipf(α, w) distribution, this limit is $1/(\alpha - 1)$; see Appendix A.2 for the

proof. Consequently, the IGP distribution that approximates the right tail of the Zipf(α, w) distribution has $\xi = 1/(\alpha - 1)$. This mirrors the result for the continuous counterpart (Pareto distribution) established in, for example, Voitalov et al. (2019).

Given these two results, it follows that the IGP approximation to the right tail of the ZP(α, θ, w) distribution has $\xi = 0$ when $\theta \in (0, 1)$, and $\xi = 1/(\alpha - 1)$ when $\theta = 1$. This means the ZP distribution is unlikely to fit both the body and the right tail of a discrete data set well, as its implied tail index, denoted by ξ_{ZP} , can take only one of the two values. Using the IGP distribution for the right tail better captures the tail behaviour.

3 Mixture model

3.1 Threshold selection

To motivate our model, we begin with an approach to threshold selection when applying the GP distribution to continuous data. As mentioned in Section 2.2, the GP distribution is appropriate for observations above a pre-selected high threshold u , regardless of the underlying generating distribution. An example of threshold selection for network and/or power-law-related data is Voitalov et al. (2019). A more principled approach (Coles, 2001) is to first fit the GP distribution over different values of u , assess the plots of the maximum likelihood estimates (MLEs) of σ and ξ , and choose the threshold above which the parameter estimates look stable, after considering their uncertainty. Whilst useful, this still requires a subjective decision, and the uncertainty in u is not accounted for. Murphy et al. (2023) proposed an objective procedure to select a single threshold that minimises the difference between the empirical and fitted Q-Q plots. However, the uncertainty around the single threshold is not considered in these approaches.

A solution to threshold selection where its uncertainty is considered is to use an extreme value mixture distributions; see Scarrott and MacDonald (2012) for a comprehensive review of continuous mixture distributions, with numerical routines for fitting these distributions available in the R package `evmix` (Hu and Scarrott, 2018). An extreme value mixture distribution combines a GP distribution for the observations above u , and another distribution H for those below u , while allowing u to vary as a parameter and therefore to be estimated. Specifically, assuming Y is a positive continuous random variable with PDF h and CDF H , we consider another positive random variable Z which equals Y if $Z \leq u$, where u is a positive threshold. Conditional on $Z > u$, Z follows the $\text{GP}(\phi_u, \sigma, \xi)$ distribution. The PDF of this mixture distribution for Z is

$$f(z) = \begin{cases} (1 - \phi_u) \times \frac{h(z)}{H(u)}, & z \leq u, \\ \phi_u \times g_u(z), & z > u, \end{cases} \quad (5)$$

where h and g_u are referred to as the bulk and tail distributions, respectively. Both ϕ_u and $H(u)$ are required for f to be a proper density i.e. to integrate to 1.

3.2 Discrete extreme value mixture distributions

To adapt the extreme value mixture model for integer-valued data, the discrete counterpart to the GP distribution, the IGP distribution, should be used for the tail. Specification of the bulk distribution is less stringent. In the case of degree distributions, the power law, or a generalisation thereof, is usually both appropriate and desirable for the bulk of a degree distribution, and hence the natural candidate is the TZP distribution.

Mathematically, a random variable X follows the $\text{TZP-IGP}(\alpha, \theta_{\text{mix}}, u, \sigma, \xi)$ distribution

if its PMF is

$$p_2(x) = \begin{cases} (1 - \phi_u)p_{\text{TZP}}(x; \alpha_{\text{mix}}, \theta_{\text{mix}}, u, w), & x = w + 1, w + 2, \dots, u, \\ \phi_u [G_u(x; \sigma, \xi) - G_u(x - 1; \sigma, \xi)], & x = u + 1, u + 2, \dots \end{cases} \quad (6)$$

where $p_{\text{TZP}}(x)$ is the PMF of the TZP distribution in Equation 3, and

$$G_u(z; \sigma, \xi) = 1 - \left[1 + \frac{\xi(z - u)}{\sigma + \xi u} \right]_+^{-1/\xi}$$

is the CDF of the GP distribution, derived from integrating $g_u(z)$ in Equation 4 under the current parameterisation. The “2” in $p_2(x)$ stands for 2 components in this mixture distribution, which is a weighted sum of the TZP and IGP distributions with weights $1 - \phi_u$ and ϕ_u , respectively.

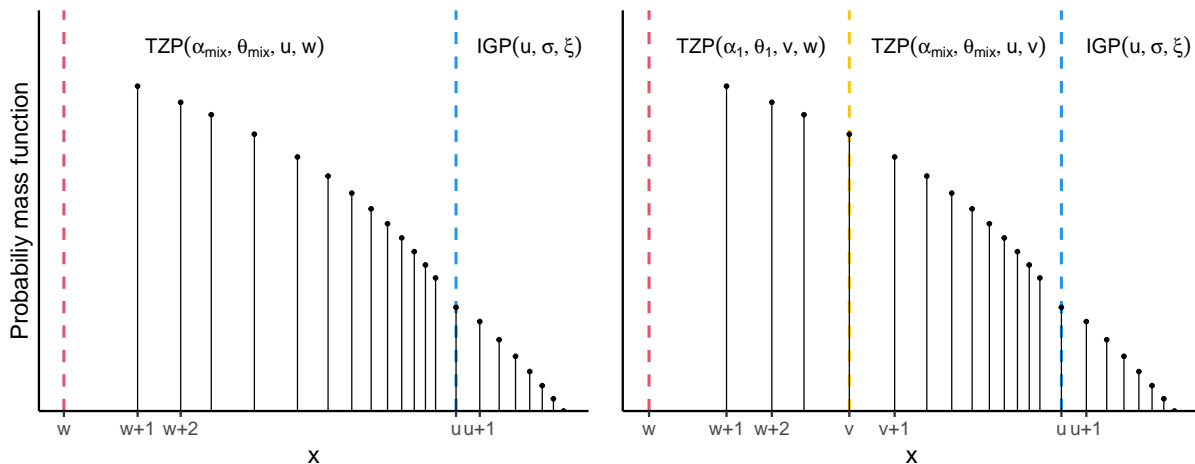


Figure 3: Schematic representations of the 2-component (left) and 3-component (right) mixture distributions. Spacing between consecutive integers is not drawn to scale.

One point regarding the parametrisation is that u is used here instead of ϕ_u , as in Section 2.2. There is no issue with this, since there is a one-to-one relationship between the two parameters, with ϕ_u the empirical proportion of exceedances once u is given; see Section 3.3. Consequently, either ϕ_u or u , but not both, can be a free parameter.

The number of components in the finite mixture modelling framework is not limited to 2; indeed a TZP-IGP mixture might still be inadequate. We also present the 3-component

counterpart, where two different TZP distributions are used for $\{w + 1, w + 2, \dots, v\}$ and $\{v + 1, v + 2, \dots, u\}$, with v an additional lower threshold. Mathematically, a random variable X follows the TZP-TZP-IGP($\alpha_1, \theta_1, v, \alpha_{\text{mix}}, \theta_{\text{mix}}, u, \sigma, \xi$) distribution if its PMF is

$$\begin{aligned}
p_3(x) &= \phi_1 \mathbb{I}_{\{x \in \{w+1, \dots, v\}\}} p_{\text{TZP}}(x; \alpha_1, \theta_1, v, w) \\
&+ (1 - \phi_1 - \phi_u) \mathbb{I}_{\{x \in \{v+1, \dots, u\}\}} p_{\text{TZP}}(x; \alpha_{\text{mix}}, \theta_{\text{mix}}, u, v) \\
&+ \phi_u \mathbb{I}_{\{x \in \{u+1, \dots\}\}} [G_u(x; \sigma, \xi) - G_u(x - 1; \sigma, \xi)], \quad x = w + 1, w + 2, \dots
\end{aligned} \tag{7}$$

where $\mathbb{I}_{\{A\}}$ is the indicator function, taking value 1 if event A occurs and 0 otherwise. We call the 2-component and 3-component mixtures the TZP-IGP and TZP-TZP-IGP distributions, respectively; their schematic representations are in Figure 3. The former can be viewed as a disjoint union of the TP-IGP (when $\theta_{\text{mix}} \in (0, 1)$) and ZM-IGP (when $\theta_{\text{mix}} = 1$) distributions. Similarly, the latter is a disjoint union of the TZP-TP-IGP and TZP-ZM-IGP distributions.

Finally, we define the *implied* tail index $\xi_{\text{mix}} = \mathbb{I}_{\{\theta_{\text{mix}}=1\}}/(\alpha_{\text{mix}} - 1)$, as discussed in Section 2.3, *had the bulk TZP distribution been extended to encompass the right tail* and provided that $\alpha_{\text{mix}} > 1$. Its difference with ξ and ξ_{ZP} , defined in Section 2.3, will be key when showing the adequacy of the mixture distributions in Section 5.

3.3 Likelihood

It is common for degree distributions to have multiple observations with the same integer value. For a sample of size n , and m unique values (x_1, x_2, \dots, x_m) with counts (c_1, c_2, \dots, c_m) such that $\sum_{i=1}^m c_i = n$, the likelihood for the 2-component mixture distribu-

tion of Equation 6 is

$$L_2(\boldsymbol{\eta}_2|\mathbf{x}, \mathbf{c}) = \prod_{i=1}^m p_2(x_i)^{c_i} = (1 - \phi_u)^{n-n_u} \phi_u^{n_u} \times \prod_{i:w < x_i \leq u} p_{\text{TZP}}(x_i; \alpha_{\text{mix}}, \theta_{\text{mix}}, u, w)^{c_i} \quad (8)$$

$$\times \prod_{i:x_i > u} [G_u(x_i; \sigma, \xi) - G_u(x_i - 1; \sigma, \xi)]^{c_i},$$

where $\boldsymbol{\eta}_2 := (w, \alpha_{\text{mix}}, \theta_{\text{mix}}, u, \sigma, \xi)$ is the parameter vector, and $n_u = \sum_{i=1}^m c_i \mathbb{I}_{\{x_i > u\}}$ is the number of exceedances of u . The parameters $(\alpha_{\text{mix}}, \theta_{\text{mix}})$ and (σ, ξ) describe the bulk and tail distributions, respectively. By definition of the model, for a given data set and for any given value of u , ϕ_u is completely determined. The converse is also true. We propose taking u to be the free parameter, in which case ϕ_u can be estimated by its MLE n_u/n , i.e. the empirical proportion of exceedances, due to the factorisation of likelihood (8). This echoes the point on the relationship between u and ϕ_u in Section 3.2, and is standard practice in extreme value methods (Coles, 2001, for example). The likelihood for the 3-component mixture follows directly:

$$L_3(\boldsymbol{\eta}_3|\mathbf{x}, \mathbf{c}) = \phi_1^{n_1} \phi_2^{n-n_1-n_u} \phi_u^{n_u} \times \prod_{i:w < x_i \leq v} p_{\text{TZP}}(x_i; \alpha_1, \theta_1, v, w)^{c_i}$$

$$\times \prod_{i:v < x_i \leq u} p_{\text{TZP}}(x_i; \alpha_{\text{mix}}, \theta_{\text{mix}}, u, v)^{c_i} \prod_{i:x_i > u} [G_u(x_i; \sigma, \xi) - G_u(x_i - 1; \sigma, \xi)]^{c_i}, \quad (9)$$

where $\boldsymbol{\eta}_3 := (w, \alpha_1, \theta_1, v, \alpha_{\text{mix}}, \theta_{\text{mix}}, u, \sigma, \xi)$. Similar to the 2-component mixture, when evaluating $L_3(\boldsymbol{\eta}_3|\mathbf{x}, \mathbf{c})$, ϕ_1 and ϕ_u are replaced by n_1/n and n_u/n , respectively, where $n_1 = \sum_{i=1}^m c_i \mathbb{I}_{\{w < x_i \leq v\}}$.

4 Inference

We now outline the Bayesian inference framework, with emphasis on the model selection in order to decide whether the body of the degree distribution follows the power law. The

2-component mixture is used to illustrate the main steps of carrying out inference; similar steps apply to the 3-component mixture.

We choose w to be pre-determined instead of inferred, as we would like it to be as low as possible to utilise as much data as possible in the inference. While 0 is a natural choice, we instead take $w = 1$ for consistency across data sets which do not have 1's. Furthermore, for data sets that do have 1's, including them usually worsens the fit. Consequently, only the 1's must be discarded before fitting the mixture distribution.

4.1 Profile likelihood of u

The frequentist framework provides a starting point of statistical inference. For fixed u , the MLEs of α_{mix} and θ_{mix} can be obtained by maximising $\prod_{i:x_i \leq u} p_{\text{TZP}}(x_i; \alpha_{\text{mix}}, \theta_{\text{mix}}, u, w)^{c_i}$. Similarly, the MLEs of σ and ξ can be obtained by maximising $\prod_{i:x_i > u} [G_u(x_i; \sigma, \xi) - G_u(x_i - 1; \sigma, \xi)]^{c_i}$. Substituting the MLEs into Equation 8 gives the profile log-likelihood, and hence the MLE, of u . While this approach is computationally efficient, the uncertainty around u cannot be quantified by using the asymptotic behaviour of the MLE, as the required regularity conditions do not hold – the profile log-likelihood usually has multiple discontinuities and is therefore not differentiable. Furthermore, the profile log-likelihood tends to be quite flat, making it difficult to separate the tail from the bulk.

In contrast, the Bayesian approach allows us to obtain the full joint posterior of $\boldsymbol{\eta}_2$ and quantify the uncertainty around u via its marginal posterior. However, as u is discrete, sampling from its posterior using all the values in the support is inefficient. Instead, we propose using its profile log-likelihood to identify the set of “most probable” values prior to applying the Bayesian inference algorithm. See Appendix B.1 for further details.

4.2 Priors

We assign independent and relatively uninformative priors as follows: $\alpha_{\text{mix}} \sim \text{N}(0, 10^2)$, $\theta_{\text{mix}} \sim \text{Beta}(1, 1)$, $\sigma \sim \text{Gamma}(1, 0.01)$, and $\xi \sim \text{N}(0, 10^2)$, where $\text{N}(m, s^2)$ denotes the Gaussian distribution with expectation m and variance s^2 , $\text{Beta}(a, b)$ the beta distribution with its common parametrisation, and $\text{Gamma}(a, b)$ the Gamma distribution with shape parameter a and rate parameter b .

The support of u varies between data sets and it seems sensible to specify the prior for u indirectly via a prior for ϕ_u . However, as previously noted, the nature of degree distributions is such that a few small values take up a large proportion of the data. Further, a prior on $(0, 1)$ is equally undesirable as ϕ_u must be sufficiently small for the IGP distribution to approximate the tail well. Consequently, we consider the proportion of *unique* exceedances of u , denoted by ψ_u . Similar to ϕ_u , once u is known, its empirical value can be computed as $\sum_{i=1}^m \mathbb{I}_{\{x_i > u\}}/m$. We assume $\psi_u \sim \text{U}(0.001, 0.9)$ *a priori* and independent of the priors for other parameters, where $\text{U}(a, b)$ denotes the continuous uniform distribution with support $[a, b]$.

4.3 Markov chain Monte Carlo and model selection

According to Bayes theorem, $\pi(\boldsymbol{\eta}_2|\mathbf{x}, \mathbf{c}) = \frac{\pi(\boldsymbol{\eta}_2, \mathbf{x}, \mathbf{c})}{\pi(\mathbf{x}, \mathbf{c})} \propto \pi(\boldsymbol{\eta}_2, \mathbf{x}, \mathbf{c}) = L_2(\boldsymbol{\eta}_2|\mathbf{x}, \mathbf{c}) \pi(\boldsymbol{\eta}_2)$, where $\pi(\mathbf{x}, \mathbf{c})$ is the marginal likelihood and $\pi(\boldsymbol{\eta}_2)$ is the joint prior for $\boldsymbol{\eta}_2$. Essentially, the posterior for $\boldsymbol{\eta}_2$ is, up to a proportionality constant, the product of the likelihood and the prior for $\boldsymbol{\eta}_2$. As the proportionality constant is computationally intractable, the joint posterior can be approximated using a Markov chain Monte Carlo (MCMC) algorithm; see Appendix B.2 for specific details of the algorithm used in Section 5.

As discussed in Section 1, in the context of degree distributions, it is of interest to

ascertain whether the distribution follows the power law, even if only in the body of the distribution. Under the mixture distribution, a power law in the body is implied when $\theta_{\text{mix}} = 1$, and therefore a simple test would be to look at the (empirical) proportion of 1's in the marginal posterior distribution for θ_{mix} . However, as θ_{mix} is continuous, there is a zero probability of drawing θ_{mix} exactly equal to 1 in the MCMC samples. Nor is it sufficient to consider θ_{mix} in a small interval close to 1 since it was shown in Section 2.3 that $\theta_{\text{mix}} \in (0, 1)$ and $\theta_{\text{mix}} = 1$ imply very different tail behaviours.

We resolve this issue by viewing the test between $\theta_{\text{mix}} \in (0, 1)$ and $\theta_{\text{mix}} = 1$ as a model selection between the TP-IGP and ZM-IGP distributions. The resulting model selection procedure requires neither that the models are nested nor that asymptotic results hold, and is embedded naturally in the MCMC algorithm (Carlin and Chib, 1995).

We first define the model variable $M := \mathbb{I}_{\{\theta_{\text{mix}}=1\}}$, and assign (non-zero) prior probabilities $\pi(M = 1)$ and $\pi(M = 0) := 1 - \pi(M = 1)$ to the two models. Next, at the end of each iteration of the MCMC algorithm, using the current value of $\boldsymbol{\eta}_2$, we set the current value of M to 1 and 0 with probabilities $\pi(M = 1|\boldsymbol{\eta}_2, \boldsymbol{x}, \boldsymbol{c})$ and $\pi(M = 0|\boldsymbol{\eta}_2, \boldsymbol{x}, \boldsymbol{c})$, respectively. The proportions of 1's and 0's in the MCMC output are the estimates of the posterior probabilities of $\theta_{\text{mix}} = 1$ and $\theta_{\text{mix}} \in (0, 1)$, denoted by $\hat{\pi}(M = 1|\boldsymbol{x}, \boldsymbol{c})$ and $\hat{\pi}(M = 0|\boldsymbol{x}, \boldsymbol{c})$, respectively. We then compute the Bayes factor $B_{10} = \frac{\hat{\pi}(M = 1|\boldsymbol{x}, \boldsymbol{c})}{\hat{\pi}(M = 0|\boldsymbol{x}, \boldsymbol{c})} \bigg/ \frac{\pi(M = 1)}{\pi(M = 0)}$; this is the evidence of “the body of the data follows the power law” against “the body of the data does not follow the power law”.

For the 3-component mixture, model selection takes place between TZP-TP-IGP ($\theta_{\text{mix}} \in (0, 1)$) and TZP-ZM-IGP ($\theta_{\text{mix}} = 1$), as whether the body between v and u follows the power law is of more interest than whether the left tail below v follows the power law.

5 Application

In this section, we show the goodness-of-fit of the proposed distributions to a number of data sets and assess the suitability of the power law in each case. Of the data sets, most are undirected, unweighted networks from the KONECT database, omitting those with clear, and potentially artificial, jumps in the degree distributions. Also included are four data sets available in the R package `powerLaw` (Gillespie, 2015), including two sets of casualty numbers in armed conflicts (“us-americans”, “native-americans”) (Bohorquez et al., 2009, Friedman, 2015, Gillespie, 2017), and two sets of word frequencies (“swiss-prot”, “moby-dick”) investigated by Bell et al. (2012) and Newman (2005), respectively. Whilst not degree distributions, whether they follow the power law has been investigated by these references. Lastly, we also consider the directed dependencies of R packages, available in the package `crandep` (Lee, 2023). As mentioned in Section 4, where a data set contains 1’s these have been removed.

For each data set, we first fit the 2-component mixture with the 3-component mixture fitted only if the former appears visually inadequate. We do not formally test between the mixtures as this is very computationally expensive. For comparison, we also fit the ZP distribution across the whole domain of the degree distribution, simultaneously carrying out model selection between the polylog ($\theta_{ZP} \in (0, 1)$) and Zipf ($\theta_{ZP} = 1$) distributions. Again, the comparison between the best mixture and the selected ZP distribution is conducted visually, since formal testing would add considerable computational complexity.

Figure 4 shows the empirical and model-based survival plots; the survival plot was chosen over the frequency plot as it magnifies the right tail. For each data set, we only show the mixture fit by either the (TZP)-ZM-IGP or (TZP)-TZP-IGP distribution based on model selection; similarly for the ZP fit by either the Zipf or polylog distribution. Due

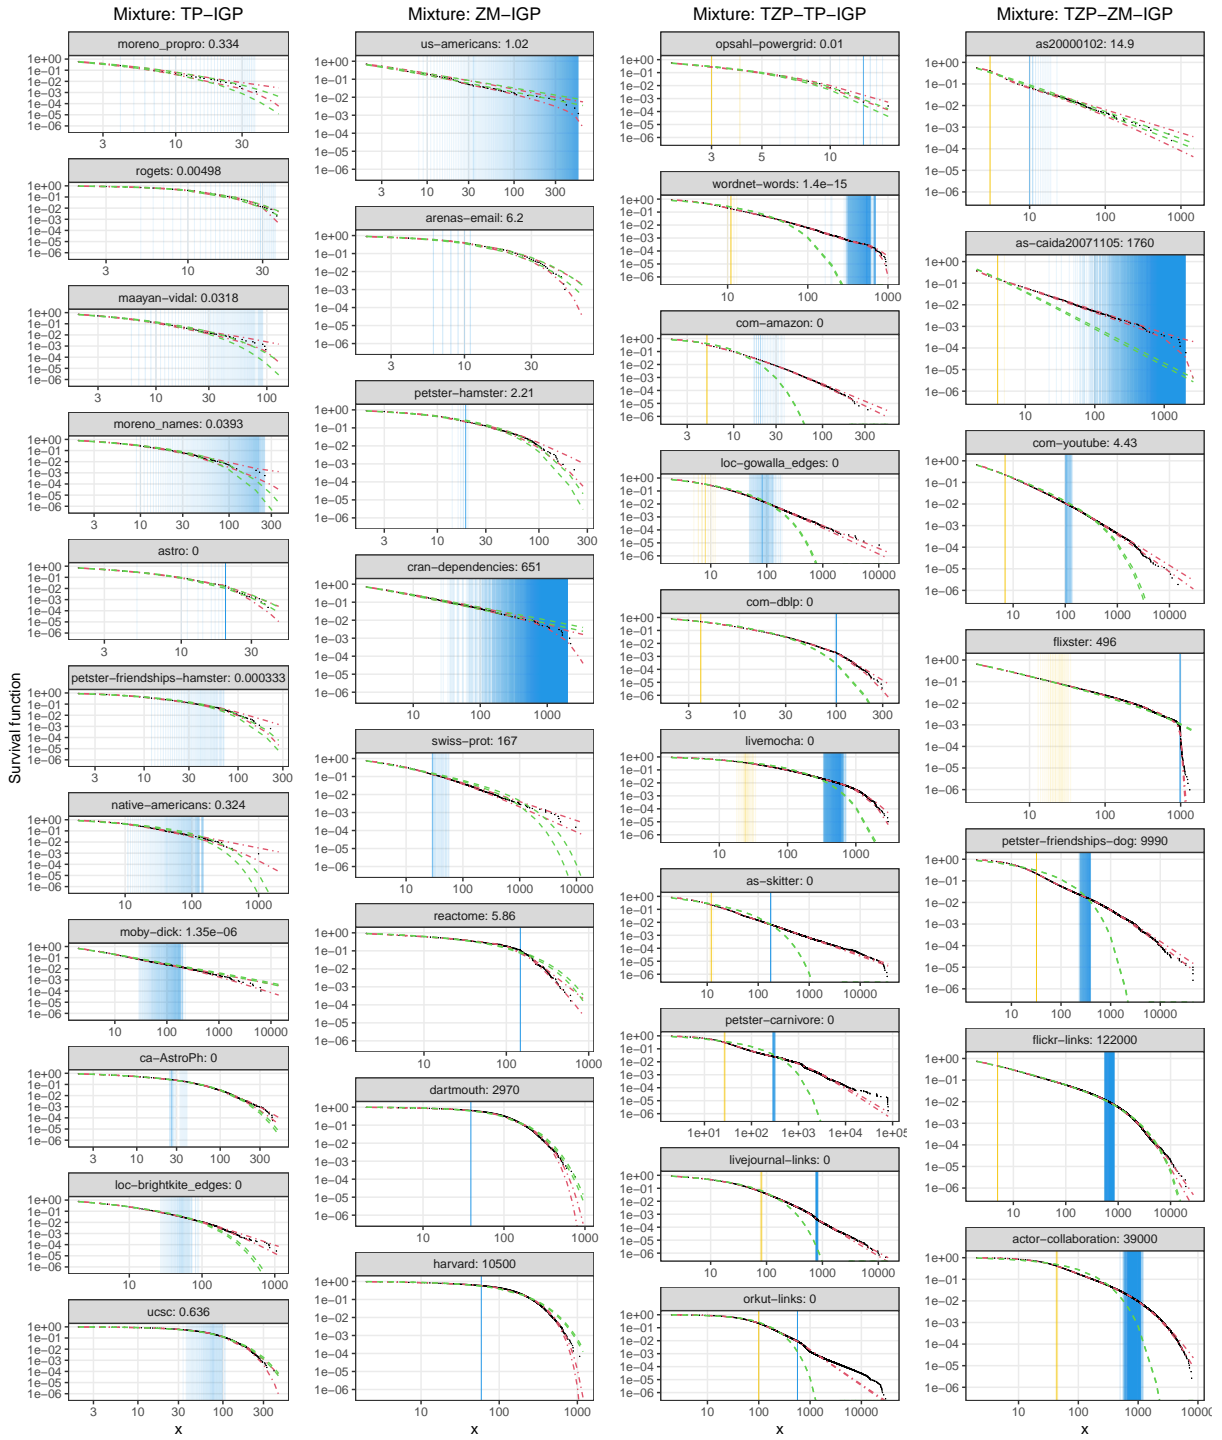


Figure 4: Survival function with credible intervals for the selected mixture (red) and ZP (green) distributions. The blue and yellow bands represent the posterior of u and v , respectively, for the mixture distribution fits. The number beside the data set name is the Bayes factor (to 3 s.f.) for $\theta_{\text{mix}} = 1$ relative to $\theta_{\text{mix}} \in (0, 1)$.

to the large sample sizes, the 95% credible intervals of the fitted survival function by the selected mixture distribution (red) are narrow for all data sets, but they do show good fits to the whole of the data. The counterparts for the fitted ZP distribution (green) show a much poorer fit in the majority of cases. This is because the curvature of the survival function leads to the polylog distribution being selected; since this implies $\xi_{\text{ZP}} = 0$, the right tail is frequently underestimated.

Figure 4 also gives insights into the uncertainty of u and, where applicable, v . The threshold u seems usually to be located to the right of the region where the power law is applicable, and the posterior support for u is a subset of the data range, demonstrating that the inference procedure for the mixture distributions locates a division between the body and the tail regions. As expected, the threshold uncertainty tends to decrease with increasing sample size, suggesting that deviations from the body are easier to identify with larger networks.

The boxplots of ξ and the implied tail indices ξ_{ZP} and ξ_{mix} in Figure 5 reinforce the need for a mixture distribution, since for most data sets, ξ of the mixture distribution is non-zero while the equivalent value implied by a ZP fit (ξ_{ZP}) is 0. Further, for the mixture fits in which the TP distribution is the bulk, ξ is markedly different to and generally larger than $\xi_{\text{mix}} = 0$. On the other hand, for cases in which the ZM distribution is the bulk, ξ is generally smaller than ξ_{mix} . Consequently, the IGP distribution captures the right tail-heaviness that $(\theta_{\text{ZP}}, \alpha_{\text{ZP}})$ in the ZP distribution cannot.

According to the definition in Section 3.2, the implied tail index ξ_{mix} is meaningful only when $\alpha_{\text{mix}} > 1$. As the sampled values are all smaller than 1 for “petster-hamster”, “dartmouth” and “harvard”, the boxplot of ξ_{mix} is therefore missing in Figure 5 for these three data sets. Had the Zipf distribution, instead of the ZM-IGP distribution, been fitted




Figure 5: Boxplot of ξ and the tail indices implied by the selected mixture and ZP distributions. The number beside the data set name is the Bayes factor (to 3 s.f.) for $\theta_{mix} = 1$ relative to $\theta_{mix} \in (0, 1)$. In the first and third columns, $\xi_{mix} = 0$; in the other two columns, some boxes for ξ_{mix} are missing as all sampled values of α_{mix} are smaller than 1.

to the whole of the data, an α smaller than 1 would not be possible.

5.1 Adequacy of the power law in the body

For each data set, we ask two questions regarding the suitability of the power law. Firstly, does the power law actually hold in the body? Equivalently, is the (TZP-)ZM-IGP distribution adequate relative to the (TZP-)TP-IGP distribution? This is answered by the Bayes factor B_{10} described in Section 4.3, in the label following the name of the data in Figure 4. According to Kass and Raftery (1995), $B_{10} > 20$ is considered as strong evidence for $\theta_{\text{mix}} = 1$ against $\theta_{\text{mix}} \in (0, 1)$. Therefore, the power law can be safely said to hold for “cran-dependencies”, “as-caida20071105”, “swiss-prot”, “dartmouth”, “harvard”, “flixster”, “petster-friendships-dog”, “flickr-links”, and “actor-collaboration”. On the other hand, “moby-dick”, “wordnet-words”, “com-amazon”, and “petster-carnivore” could be deemed to follow the power law from the log-log plot, however their Bayes factors are all effectively zero, suggesting otherwise. There is no contradiction here – although the posterior for θ_{mix} for these four data sets is highly concentrated close to 1, proximity to 1 does not translate to a high Bayes factor since $\theta_{\text{mix}} \in (0, 1)$ and $\theta_{\text{mix}} = 1$ imply different right tail behaviours.

Secondly, for the data sets for which the power law holds in the body, is it plausible that it holds for the whole range of values (for the 2-component mixture) or the data above v (for the 3-component mixture)? Among those identified above to (partially) follow the power law, “cran-dependencies” and “as-caida20071105” are the most likely candidates since both data sets have a large amount of uncertainty of u and the posterior has an upper end point close to the maximum data value. For “as-caida20071105”, the green credible intervals in Figure 4 indicates a poor fit of the selected Zipf distribution to the whole of the

data; fitting to the data above v would result in a better fit that is closer to the mixture counterpart.

Our findings strengthen the need for the mixture distributions, and highlight the consequence of fitting the incorrect distribution. The strong evidence for the partial power law for the data sets identified in this section means that the PA model could be the network generative model *up to a point*. Applying the ZP distribution with model selection to the whole of the data often leads to the polylog distribution being selected, thus overlooking the power law in the body and the potential of partial scale-free growth of the network.

6 Discussion

In this paper, we proposed a modelling and inference framework for the application of discrete extreme value mixture distributions in situations in which the power law would otherwise automatically be applied. The mixture distributions are shown to provide good fits to both the body and the tail of the data for numerous examples, whilst also quantifying the threshold uncertainty and without subjectively discarding most data points as in a traditional extreme value analysis. For most of the data sets analysed, the popular Zipf distribution is insufficient to describe the right tail of the data, which is usually lighter than would be expected from the decay of the body of the data, and nor is its generalisation – the ZP distribution – adequate to describe the whole of the data. The mixture distribution capitalises on the usefulness of the ZP distribution by utilising its truncated version – the TZP distribution – for the body of the data, while capturing the right tail using the IGP distribution.

Bayesian model selection enables us to formally determine whether the power law is adequate for the body of the data fitted by the TZP distribution. It is possible to expand the

model selection to between the ZP and mixture distributions, essentially testing if $u = \infty$ is adequate, or between the IGP distribution and the (2-component) mixture, essentially testing if $u = 1$ is adequate. However, we deem neither model selection necessary, as it is clear from the applications that the mixture distribution outperforms the ZP distribution, and as for most data sets u is comfortably away from the lower and upper end-points of the distribution, indicating the usefulness of the mixture distributions. On the other hand, if model selection between all distributions (polylog, Zipf, TP-IGP, ZM-IGP, TZP-TP-IGP, TZP-ZM-IGP) is necessary, a systematic approach would be to compare the marginal likelihood for each distribution, which can be computed within the MCMC algorithm.

Lastly, we return to the connection between the degree distribution and network generating models. Further investigations on this are beyond the scope of this paper, however the results in Section 5 suggest that the PA model is not a realistic generating mechanism for many real-world examples of degree distributions. Given the larger number of networks for which the mixture distribution best describes the degree distribution, we feel that it is of great importance to investigate whether the PA or similar network generative models could be easily modified to result in the desired degree distributions. However if the growth according to the preferential attachment model is to be studied in real networks, the two data sets identified at the end of Section 5 (“cran-dependencies” and “as-caida20071105”) would be the most suitable candidates for such an investigation.

References

Albert, R., Jeong, H. and Barabási, A.-L. (1999), ‘Internet: Diameter of the world-wide web’, *Nature* **401**, 130–131.

- Arroyo-Machado, W., Torres-Salinas, D., Herrera-Viedma, E. and Romero-Frías, E. (2020), ‘Science through Wikipedia: A novel representation of open knowledge through co-citation networks’, *PLOS One* **15**(2), e0228713.
- Barabási, A.-L. and Albert, R. (1999), ‘Emergence of scaling in random networks’, *Science* **286**(5439), 509–512.
- Bell, M. J., Gillespie, C. S., Swan, D. and Lord, P. (2012), ‘An approach to describing and analysing bulk biological annotation quality: A case study using UniProtKB’, *Bioinformatics* **28**(18), i562–i568.
- Bhamidi, S., Steele, J. M. and Zaman, T. (2015), ‘Twitter event networks and the Superstar model’, *The Annals of Applied Probability* **25**(5), 2462–2502.
- Bohorquez, J. C., Gourley, S., Dixon, A. R., Spagat, M. and Johnson, N. F. (2009), ‘Common ecology quantifies human insurgency’, *Nature* **462**(7275), 911–914.
- Broido, A. D. and Clauset, A. (2019), ‘Scale-free networks are rare’, *Nature Communications* **10**(1017).
- Carlin, B. P. and Chib, S. (1995), ‘Bayesian model choice via Markov chain Monte Carlo methods’, *Journal of the Royal Statistical Society: Series B (Methodological)* **57**(3), 473–484.
- Clauset, A., Shalizi, C. R. and Newman, M. E. J. (2009), ‘Power-law distributions in empirical data’, *SIAM Review* **51**(4), 661–703.
- Coles, S. (2001), *An Introduction to Statistical Modeling of Extreme Values*, Springer Series in Statistics, Springer.

- Faloutsos, M., Faloutsos, P. and Faloutsos, C. (1999), On power-law relationships of the internet topology, *in* ‘Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication’, SIGCOMM ’99, ACM, New York, NY, USA, pp. 251–262.
- Friedman, J. A. (2015), ‘Using power laws to estimate conflict size’, *Journal of Conflict Resolution* **59**, 1216–1241.
- Gillespie, C. S. (2015), ‘Fitting heavy tailed distributions: the powerLaw package’, *Journal of Statistical Software* **64**(2), 1–16.
- Gillespie, C. S. (2017), ‘Estimating the number of casualties in the American Indian war: a Bayesian analysis using the power law distribution’, *Annals of Applied Statistics* **11**(4), 2357–2374.
- Hitz, A., Davis, R. and Samorodnitsky, G. (2017), ‘Discrete extremes’, *ArXiv e-prints* . arXiv:1707.05033.
- Hu, Y. and Scarrott, C. (2018), ‘evmix: An R package for extreme value mixture modeling, threshold estimation and boundary corrected kernel density estimation’, *Journal of Statistical Software* **84**(5), 1–27.
- Johnson, N. L., Kemp, A. W. and Kotz, S. (2005), *Univariate Discrete Distributions*, Wiley Series in Probability and Statistics, 3rd edn, Wiley.
- Johnston, J. and Andersen, T. (2022), ‘Random processes with high variance produce scale free networks’, *Physica A* **604**.
- Jung, H. and Phoa, F. K. H. (2021), ‘A mixture model of truncated zeta distributions with applications to scientific collaboration networks’, *Entropy* **23**(5).

- Kass, R. E. and Raftery, A. E. (1995), ‘Bayes factors’, *Journal of the American Statistical Association* **90**(430), 773–795.
- Kunegis, J. (2013), KONECT - The Koblenz Network Collection, *in* ‘Proceedings of the International Conference on World Wide Web Companion’, pp. 1343–1350.
- Lee, C. (2023), *crandep: Network Analysis of Dependencies of CRAN Packages*. R package version 0.3.2.
- Lee, J. and Oh, J. C. (2014), Estimating the degrees of neighboring nodes in online social networks, *in* ‘PRIMA 2014: Principles and Practice of Multi-Agent Systems’, Lecture Notes in Artificial Intelligence, Springer.
- Mathews, P., Mitchell, L., Nguyen, G. and Bean, N. (2017), The nature and origin of heavy tails in retweet activity, *in* ‘Proceedings of the 26th International Conference on World Wide Web Companion’, WWW ’17 Companion, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, pp. 1493–1498.
- Murphy, C., Tawn, J. A. and Varty, Z. (2023), ‘Automated threshold selection and associated inference uncertainty for univariate extremes’, *ArXiv e-prints* . arXiv:2310.17999.
- Newman, M. E. J. (2005), ‘Power laws, Pareto distributions and Zipf’s law’, *Contemporary Physics* **46**(5), 323–351.
- Pickands, J. (1975), ‘Statistical inference using extreme order statistics’, *The Annals of Statistics* **3**(1), 119–131.
- Price, D. (1976), ‘A general theory of bibliometric and other cumulative advantage pro-

- cesses', *Journal of the Association for Information Science and Technology* **27**(5), 292–306.
- Prieto, F., Gómez-Déniz, E. and Sarabia, J. M. (2014), 'Modelling road accident blackspots data with the discrete generalized Pareto distribution', *Accident Analysis and Prevention* **71**, 38–49.
- Rohrbeck, C., Eastoe, E. F., Frigessi, A. and Tawn, J. A. (2018), 'Extreme value modelling of water-related insurance claims', *The Annals of Applied Statistics* **12**(1), 246–282.
- Scarrott, C. and MacDonald, A. (2012), 'A review of extreme value threshold estimation and uncertainty quantification', *REVSTAT - Statistical Journal* **10**(1), 33–60.
- Sheridan, P. and Onodera, T. (2018), 'A preferential attachment paradox: How preferential attachment combines with growth to produce networks with log-normal in-degree distributions', *Scientific Reports* **8**(2811).
- Shimura, T. (2012), 'Discretization of distributions in the maximum domain of attraction', *Extremes* **15**, 299–317.
- Stumpf, M. P. H. and Porter, M. A. (2012), 'Critical truths about power laws', *Science* **335**(6069), 665–666.
- Stumpf, M. P. H., Wiuf, C. and May, R. M. (2005), 'Subnets of scale-free networks are not scale-free: sampling properties of networks', *Proceedings of the National Academy of Sciences* **102**(12), 4221–4224.
- Thelwall, M. and Wilson, P. (2014), 'Regression for citation data: An evaluation of different methods', *Journal of Infometrics* **8**, 963–971.

- Valero, J., Pérez-Casany, M. and Duarte-López, A. (2022), ‘The Zipf-Polylog distribution: Modeling human interactions through social networks’, *Physica A* **603**.
- Varga, I. (2015), Scale-free network topologies with clustering similar to online social networks, *in* H. Takayasu, N. Ito, I. Noda and M. Takayasu, eds, ‘Proceedings of the International Conference on Social Modeling and Simulation, plus Econophysics Colloquium 2014’, Spring Proceedings in Complexity, Springer, pp. 323–333.
- Voitalov, I., van der Hoorn, P., van der Hofstad, R. and Krioukov, D. (2019), ‘Scale-free networks well done’, *Phys. Rev. Res.* **1**, 033034.
- Yule, G. U. (1925), ‘II.-A mathematical theory of evolution, based on the conclusions of Dr. J. C. Willis, F. R. S’, *Philosophical Transactions of the Royal Society of London. Series B, Containing Papers of a Biological Character* **213**(402-410), 21–87.