



BIROn - Birkbeck Institutional Research Online

Fenner, Trevor and Levene, Mark and Loizou, George (2016) A stochastic evolutionary model generating a mixture of exponential distributions. The European Physical Journal B (EPJ B) 89 (2), ISSN 1434-6036.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/14373/>

Usage Guidelines:

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html>
contact lib-eprints@bbk.ac.uk.

or alternatively

A stochastic evolutionary model generating a mixture of exponential distributions

Trevor Fenner, Mark Levene, and George Loizou
Department of Computer Science and Information Systems
Birkbeck, University of London
London WC1E 7HX, U.K.
{mark,trevor,george}@dcs.bbk.ac.uk

Abstract

Recent interest in human dynamics has stimulated the investigation of the stochastic processes that explain human behaviour in various contexts, such as mobile phone networks and social media. In this paper, we extend the stochastic urn-based model proposed in [FLL15] so that it can generate mixture models, in particular, a mixture of exponential distributions. The model is designed to capture the dynamics of survival analysis, traditionally employed in clinical trials, reliability analysis in engineering, and more recently in the analysis of large data sets recording human dynamics. The mixture modelling approach, which is relatively simple and well understood, is very effective in capturing heterogeneity in data. We provide empirical evidence for the validity of the model, using a data set of popular search engine queries collected over a period of 114 months. We show that the survival function of these queries is closely matched by the exponential mixture solution for our model.

Keywords: human dynamics, stochastic model, survival analysis, mortality rate, mixture model, exponential mixture

1 Introduction

Recent interest in complex systems, such as social networks, the world-wide-web, email networks and mobile phone networks [Bar07], has led researchers to investigate the processes that could explain the dynamics of human behaviour within these networks. Barabási [Bar05] has suggested that the bursty nature of human behaviour, for example, when measuring the inter-event response time of email communication, is a result of a decision-based queuing process. In particular, humans tend to prioritise actions, for example, when deciding which email to respond to, and therefore a priority queue model was proposed in [Bar05], leading to a heavy-tailed power-law distribution of inter-event times. The availability of large data sets, such as mobile phone records, has widened the applicability of human dynamics investigation, for example, in an attempt by Schneider et al. [SBCG13] to uncover the characteristics of daily mobility patterns. Human dynamics is not limited to the study of behaviour within communication networks, as can be seen, for example, by a recent proposal of Mitnitski et al. [MSR13], who apply a simple stochastic queuing model to the complex phenomenon of ageing, in order to illustrate how health deficits accumulate with age. Moreover, *sociophysics* [Gal08, CFL09, SC14] has similar aims and, in particular, uses concepts and methods from

statistical physics to investigate social phenomena and political behaviour. This methodology often involves the empirical investigation of collective choice dynamics, resulting in the formation of statistical laws governing, for example, how the popularity of movies emerges from social behaviour. In this context, we have recently modelled the survival of search engine queries in the top-10 list of the most popular queries that users have submitted, as an application of human dynamics [FLL15]; see [TdFA14] for attempts to predict the future popularity of online content.

Survival analysis [Law03, KK12] provides statistical methods to estimate the time until an event occurs, known as the *survival* or *failure* time. Typically, an event in a survival model is referred to as a failure, since it often has negative connotations, such as mortality or the contraction of a disease, although it could, in principle, also be positive, such as the time to return to work or to recover from a disease. In the context of email communication mentioned above, an event might be a reply to an email; in the context of popular queries, an event might be dropping out of the top-10 most popular queries. Traditional applications of survival analysis are in clinical trials [FL00], and in understanding the mechanisms in biological ageing [GG01]. The methods used in survival analysis overlap with those used in engineering for reliability life data analysis [Fin08, OK12]. Reliability analysis in engineering has many applications, for example, in manufacturing processes, and in software design and testing. However, one can envisage that survival analysis would find application in newer human dynamics scenarios in complex systems, such as those arising in social and communication networks [Bar05, CGW⁺08, MAAJ13].

We have been particularly interested in formulating *generative models* in the form of stochastic processes by which complex systems evolve and give rise to power laws or other distributions [FLL07, FLL12, FLL15]. This type of research builds on the early work of Simon [Sim55], and the more recent work of Barabási's group [AB02] and other researchers [BSV07]. In the context of human dynamics, the priority queue model mentioned above [Bar05] is a generative model characterised by a heavy-tailed distribution.

In the bigger picture, one can view the goal of such research as being similar to that of *social mechanisms* [HS98], which looks into the processes, or mechanisms, that can explain observed social phenomena. Using an example given in [Sch98], the growth in the sales of a book can be explained by the well-known logistic growth model [TW02], and more recently we have shown that the process of conference registration with an early bird deadline can be modelled by bi-logistic growth [FLL13]. Such research can also be grounded within the field of sociophysics, mentioned above, where critical phenomena are important in demonstrating how the transition to global behaviour can emerge from the interactions of many individuals. The individuals may be neurons, as in [LAE⁺12], where criticality emerges as neuron cooperation, or people's decisions, as in [PS10], where the popularity of movies emerges as collective choice behaviour.

In this work we build on the generative model proposed in [FLL15], which captures the essential dynamics of survival analysis applications. In the urn-based model presented in [FLL15], the *actors* are called *balls*, and a ball being present in *urn_i*, the *i*th urn, indicates that the actor represented by the ball has so far survived for *i* time steps. An actor could, for example, be a subject in a clinical trial, an email that has not yet been replied to, or an ongoing phone call. As a simplification, we assume that time is discrete and that, at any given time, a single ball may join the system with a fixed *birth* probability. In addition, existing balls in the system may be discarded, i.e. one or more *mortality* (or death) events may occur,

according to some specified *mortality function*, which may take a general form. It is evident that, at any given time t , we may have at most one ball in urn_i , for all $i \leq t$.

The probability that, after t steps of the stochastic process outlined above, there is a surviving ball in urn_i is known as the *survival function* [KK12]. In the context of a medical study, this could represent the probability that a patient survives for longer than a specified time. In [FLL15] we have shown how the survival function depends on the mortality function. Two important cases are: (1) when the mortality function is constant then the survival function follows an exponential distribution, and (2) when the mortality function follows a power-law distribution then the survival function follows a Weibull distribution [Rin09]; the latter is often applied to the modelling of *life data* (also known as survival data, time to event data, or time to failure data) [KK12, OK12].

It is interesting to observe that the distribution of the survival function has two parameters, i and t , cf. [FLL12], whereas most previously studied generative stochastic models [AB02], including those in our previous work, for example [FLL07], result in steady state distributions that are asymptotic in t to a distribution with a single parameter i .

Exponential mixture models [MP00] have been exploited in survival analysis in a variety of applications, as mentioned in Section 2 below. Moreover, it has been shown that the survival function of an exponential mixture model has a decreasing failure rate [BJ97, Fin09], which was first observed in [Pro63] in the context of air-conditioning systems for a fleet of airplanes. To generate exponential mixtures, in Section 3, we extend the urn-based stochastic process of [FLL15] by adding a first step of tossing a biased coin to decide which of two processes becomes active; although the extension is defined for a mixture of two exponentials, it extends naturally to a mixture of several exponentials. We provide empirical evidence for the validity of the model by analysing the longitudinal data set of popular search engine queries over 114 months that we used in [FLL15]. We show that the survival function of these queries closely follows the distribution generated by our model.

The rest of the paper is organised as follows. In Section 2, we review mixture models in the context of survival data analysis. In Section 3, we present our stochastic urn-based model that provides us with a mechanism to model the mixture of exponential distributions. In Section 4, we apply the model to the survival of popular search engine queries posted on Google Trends (www.google.com/trends/hottrends). Finally, in Section 5, we give our concluding remarks.

2 Mixture models for Failure Time Analysis

Lifetime analysis, also referred to as survival time analysis or failure time analysis [Law03], has a long history of exploiting mixture models [MP00], in particular exponential mixture models, in a variety of applications. Examples of areas where exponential mixtures have been employed for lifetime analysis modelling are: the reliability of system components [Pro63], network performance [FW98], actuarial losses [Kea99], reliability engineering [Fin09], and the likelihood of volcanic eruption [MD09].

The application we consider in Section 4, in the context of survival analysis, is that of predicting the duration of the popularity of search engine queries. We note that this is related to the problem of predicting the popularity of movies [PS10] or online content [TdFA14]. Mixture models are a natural tool for modelling such applications as they are very effective in

capturing heterogeneity. In particular, when considering the popularity of queries, a mixture model can differentiate between two types of queries: ones that are very popular for a short period of time and others whose popularity lasts much longer; such an observation has already been made in the context of the popularity of movies [PS10].

The *survival function* $S(\theta)$ represents the probability that an actor survives for longer than a specified time θ . Here we restrict ourselves to a mixture of two exponential survival functions, noting that this can readily be extended to a mixture of more than two such functions. The mixture rates of the distribution, p_1 and p_2 , sum to one, and the *mortality rates*, λ_1 and λ_2 , of the mixture model (also known as the *failure rates*) are constants. These define a *discrete hazard function* [BG03], i.e. the conditional probability that a failure occurs at a specified time instant given that no failure occurred previously. We therefore define the mixture survival function $S(\theta)$ as

$$S(\theta) = p_1 \exp(-\lambda_1 \theta) + p_2 \exp(-\lambda_2 \theta), \quad (1)$$

where $p_1 + p_2 = 1$.

The mortality rate of the mixture model [Fin09] is now given by

$$\lambda(\theta) = \pi_1(\theta)\lambda_1 + \pi_2(\theta)\lambda_2, \quad (2)$$

where the time-dependent probabilities π_k are defined as

$$\pi_k(\theta) = \frac{p_k \exp(-\lambda_k \theta)}{p_1 \exp(-\lambda_1 \theta) + p_2 \exp(-\lambda_2 \theta)}. \quad (3)$$

In [BJ97, Fin09] it was shown that the mortality rate $\lambda(\theta)$ of the mixture model is decreasing in θ . Moreover, if we assume that $\lambda_1 \leq \lambda_2$, then $\lambda(\theta)/\lambda_1$ converges to 1 as θ tends to infinity; this is referred to in the literature as a *decreasing failure rate*.

For comparison purposes, we will also make use of the Weibull survival function considered in [FLL15], with shape parameter ρ and scale parameter γ , which is given by

$$\exp(-\gamma \theta^{1+\rho}). \quad (4)$$

It is important to note that, when the shape parameter satisfies $-1 < \rho < 0$, then it can be expressed as a mixture of exponentials, since it is a *completely monotone function* [Jew82, FW98]; see [Fel71] for the underlying theory.

3 An Evolutionary Urn Transfer Model

We next formalise a stochastic urn model based on the survival model for human dynamics introduced in [FLL15], in order to generate the mixture of two exponential distributions. This is done by first generating a mixture distribution by tossing a coin with mixture probabilities p_1 and p_2 , where $p_1 + p_2 = 1$, to decide which process is active, and then employing a constant mortality rate parameter λ_k in order to generate an exponential distribution, as in [FLL15]. We note that this mechanism extends naturally to a mixture of more than two exponentials by deploying a multi-sided “coin”.

We assume a countable number of urns, urn_0, urn_1, \dots , where a *ball* (or *actor*) in urn_i is said to be of *age* i . Initially, all the urns are empty. At the start of the stochastic process, we

toss the coin, as described above, to decide which of two processes will be active. From then on, at any time t , $t \geq 0$, a new ball may be *born*, with probability b , and any such new ball is inserted into urn_0 ; in addition, an existing ball of age i can *die* by being discarded from urn_i , for all $i > 0$.

For a given age i and time t , we let the *mortality rate* $\mu_k(i, t)$ be the probability that a ball in urn_i dies at time t , where $k = 1$ or 2 depending on which process is active. We always require that $\mu_k(0, t) = 0$ for all t , and that $\mu_k(i, t) = \lambda_k$ for $i > 0$, where λ_k is a constant, satisfying $0 < \lambda_k < 1$.

At time t , the stochastic process then proceeds as follows in order to obtain the configuration at time $t + 1$, where $t \geq 0$.

- (i) For each i , $1 \leq i \leq t$, if urn_i is non-empty then, with probability $\mu_k(i, t) = \lambda_k$, a mortality event occurs, i.e. the ball in urn_i is discarded.
- (ii) Next, the ages of all balls remaining in the system are incremented by 1, i.e. a ball in urn_i is moved to urn_{i+1} , for each i .
- (iii) Finally, with probability b , where $0 < b < 1$, a birth event occurs, i.e. a ball is inserted into urn_0 .

We observe that, at any time t , urn_i is empty for $i > t$. Note that we are assuming that the mortality rate is constant, i.e. $\mu_k(i, t) = \lambda_k$, which leads to the survival function of an exponential distribution, as was shown in [FLL15]. However, in principle, the mortality rate could be an arbitrary function, for example, power-law mortality leads to the survival function of a Weibull distribution. Note also that we are using a single birth probability b , rather than one for each process, since, as can be seen from (7) below, b is merely a scaling factor. It follows that having different values for b for the two processes is equivalent to adjusting the mixture probabilities.

We now let $F_k(i, t) \geq 0$, with $k = 1$ or 2 , be a discrete function denoting the probability that there is a ball in urn_i at time t , given that process k is active. Initially, we set $F_k(0, 0) = b$ and $F_k(i, 0) = 0$ for all $i > 0$.

The dynamics of the model are captured by the following two equations, where we assume that $k = 1$ or 2 according to the outcome of the coin toss:

$$F_k(0, t) = b \text{ for } t \geq 0, \tag{5}$$

and

$$F_k(i + 1, t + 1) = F_k(i, t) - \mu_k(i, t)F_k(i, t) \text{ for } 0 \leq i \leq t. \tag{6}$$

We can expand (6) to obtain

$$F_k(i + 1, t + 1) = b \prod_{j=1}^i (1 - \mu_k(j, t - i + j)) = b(1 - \lambda_k)^i. \tag{7}$$

Using the approximation $\ln(1 - x) \approx -x$, which holds for small x , we see that the distribution is approximately exponential:

$$F_k(i + 1, t + 1) \approx b \exp(-\lambda_k i). \tag{8}$$

We note that, in (8), we could use the more accurate second order approximation $\ln(1 - x) \approx -x - x^2/2$.

We now define $F(i, t)$ analogously to $F_k(i, t)$ to take into account which of the two processes is active, i.e.

$$F(i, t) = p_1 F_1(i, t) + p_2 F_2(i, t).$$

This yields the initial conditions $F(0, 0) = b$, $F(i, 0) = 0$ for all $i > 0$, and also $F(0, t) = b$ for $t \geq 0$. Similarly, from (7) and (8) for $0 \leq i \leq t$, we obtain the exponential mixture:

$$\begin{aligned} F(i + 1, t + 1) &= p_1 b(1 - \lambda_1)^i + p_2 b(1 - \lambda_2)^i \\ &\approx p_1 b \exp(-\lambda_1 i) + p_2 b \exp(-\lambda_2 i). \end{aligned} \quad (9)$$

We note that, as mentioned in the beginning of this section, the model can be extended to more than two components; in this case the sums in (9) would have the appropriate number of terms.

4 Application of the Mixture Model

In the context of human dynamics, survival analysis has recently been applied to large data sets. These include the analysis of phone call durations [VAFL10], the investigation of how long Wikipedia editors remain active [ZPL12], and for predicting the likelihood of online content to become popular [LMS12].

The survival function $S(\theta)$ introduced in Section 2 is usually estimated via a step function by computing the probability that an actor survives until time θ , for $\theta = 1, 2, \dots, t$; this step function is known as the *Kaplan-Meier estimator* [KM58, KK12]. By comparing (7) with the Kaplan-Meier estimators for the survival functions [KM58, equation (2b)] of the components of the mixture model, these are seen to be analogous to $F_k(i, t)$ for $k = 1, 2$, and an actor that was born at time $t - i$. We thus obtain

$$S(i) \approx \frac{p_1 F_1(i, t) + p_2 F_2(i, t)}{b} = \frac{F(i, t)}{b}. \quad (10)$$

In theory, the survival function $S(\theta)$ does not depend on the length t of the trial. However, in practice, the Kaplan-Meier estimate will be more accurate for longer trials. Nevertheless, this estimate is more accurate when most of the actors are still present in the study, since the estimate may be inaccurate when there are only a few actors remaining [RNP⁺10].

As as a proof of concept for the exponential mixture model, we analysed the survival of queries in the top-10 Google Trends “hot searches” (www.google.com/trends/hottrends), that we used in [FLL15]. Data was collected monthly for the top-10 “hot searches” over 114 months, from January 2004 until June 2013, for six categories (together with their subcategories in each case): Business & Politics (or simply Business), Entertainment, Nature & Science (or simply Science), Shopping & Fashion (or simply Shopping), Sports, and Travel & Leisure (or simply Travel). The number of distinct queries per category over the period is shown in Table 1. It is apparent from this statistic that the top-10 queries from Shopping change the least, while those from Entertainment change the most.

In this data set, the balls are top-10 queries, a mortality event occurs when a query leaves the top-10 in a given month, and a birth event occurs when a new query joins the top-10 in

Data set	Number of queries
Business	318
Entertainment	1672
Science	150
Shopping	107
Sports	774
Travel	342
All Categories	3363

Table 1: Number of top-10 queries collected from Google Trends.

a given month (note that time is discrete and is measured in months). As we are considering an exponential mixture model with constant mortality for each component, queries leave the top-10 according to (2).

We next outline the methodology we have used to validate and evaluate the stochastic urn model presented in Section 3. We then give further details, before discussing and analysing the results.

- (I) First, to obtain estimates of p_1 , λ_1 and λ_2 , we performed nonlinear regression to the right-hand side of (9) using the Kaplan-Meier estimates computed from the raw data, for $i = 1, 2, \dots, 114$.
- (II) We used the estimates of p_1 , λ_1 and λ_2 from (I) to compute, for each i , the right-hand side of (9); we call this the *mixture data*. We then repeated the nonlinear regression using the mixture data as a quick “sanity check” that the new values of p_1 , λ_1 and λ_2 obtained were consistent with those from (I).
- (III) Next we ran simulations using the parameters p_1 , λ_1 and λ_2 from (I), and $b = 0.9$. Using the averaged values of $F(i, t)$ from the simulations, we again repeated the nonlinear regression to obtain new values for p_1 , λ_1 and λ_2 ; these were compared with those from (I) for consistency.
- (IV) We computed the D values for Kolmogorov-Smirnov tests to ascertain whether the Kaplan-Meier estimates, the mixture data and the averaged simulation data are likely to have all come from the same distribution.

We first obtained the Kaplan-Meier estimates from the raw data sets for the six individual categories and for their aggregation (All Categories). Recalling that the survival function is approximated as in (10), following (I), we used Matlab to obtain estimates for p_1 , λ_1 and λ_2 by nonlinear least-squares regression of $S(i)$ on i for fitting the right-hand side of (9), for $i = 1, 2, \dots, 114$. The estimated parameters p_1 , λ_1 and λ_2 are shown in the rows of Table 2, together with the coefficient of determination R^2 [Mot95]; the R^2 values show a very good fit for all of the categories. Nonlinear regression using the mixture data values computed from the right-hand side of (9), as in (II), yielded almost perfect fits, as expected.

To test the validity of the model, as in (III), we then carried out simulations in Matlab of the stochastic urn transfer model using the values of p_1 , λ_1 and λ_2 from Table 2. We chose

Data set	p_1	λ_1	λ_2	R^2
Business	0.4061	0.0136	0.4095	0.9893
Entertainment	0.2013	0.0405	0.4793	0.9965
Science	0.6693	0.0071	0.4391	0.9884
Shopping	0.7403	0.0091	0.4880	0.9915
Sports	0.2275	0.0215	0.4599	0.9829
Travel	0.4483	0.0090	0.4364	0.9823
All Categories	0.2600	0.0154	0.4337	0.9894

Table 2: Nonlinear least-squares regression to a mixture of exponentials using the Kaplan-Meier estimates.

the value $b = 0.9$ for all simulations, after running some sample simulations with other values of b . The value of b is not critical since, as can be seen from (9), b is merely a scaling factor. The simulations were run for 114 steps, one for each month, for each of the categories, and this was repeated 10^6 times. For each category, we then calculated the average value of $F(i, t)$ for $i = 1, 2, \dots, t$, over the 10^6 runs. Repeating the nonlinear regression using these average values gave the results shown in Table 3. Comparing Table 3 with Table 2 shows that all the values of p_1 , λ_1 and λ_2 are very close. Moreover, the very high values for R^2 confirm the validity of (9) as the solution to our model.

Data set	p_1	λ_1	λ_2	R^2
Business	0.4052	0.0135	0.4037	0.9999
Entertainment	0.1991	0.0402	0.4683	0.9996
Science	0.6686	0.0071	0.4289	0.9994
Shopping	0.7394	0.0091	0.4760	0.9996
Sports	0.2269	0.0214	0.4504	0.9997
Travel	0.4472	0.0089	0.4286	0.9997
All Categories	0.2598	0.0153	0.4260	0.9998

Table 3: Nonlinear least-squares regression to a mixture of exponentials using the simulated data with p , λ_1 and λ_2 taken from Table 2.

To compare the Kaplan-Meier estimates, the mixture data and the averaged simulation data, as in (IV), we performed three Kolmogorov-Smirnov 2-sample 2-tailed tests, as described in Section 6.6.4 of [SC88]. Taking the null hypothesis to be that the Kaplan-Meier estimates, the mixture data and the averaged simulation data all come from the same population distribution, the critical value at significance level $\alpha = 0.05$ is 0.1801 for a sample of 114 (number of months). The D values for the three pairwise tests are shown in Table 4. It can be seen that, in all cases, the values of the test statistic D are well below the critical value. Hence, we cannot reject the null hypothesis at significance level $\alpha = 0.05$. The values in the Sim-Mix column show that the distributions of the mixture data and averaged simulation data are extremely close, which is unsurprising since these are both derived directly from our model. We note that, even at significance level $\alpha = 0.10$, where the critical value is 0.1616, the null hypothesis cannot be rejected.

Data set	KM-Sim	KM-Mix	Sim-Mix
Business	0.0761	0.0720	0.0126
Entertainment	0.0640	0.0583	0.0247
Science	0.0727	0.0598	0.0339
Shopping	0.0351	0.0343	0.0289
Sports	0.1015	0.0966	0.0226
Travel	0.0841	0.0807	0.0170
All Categories	0.0856	0.0803	0.0175

Table 4: The D values for the 2-sample 2-tailed Kolmogorov-Smirnov tests.

In [FLL15], we used nonlinear least-squares regression to fit Weibull survival functions to the Kaplan-Meier estimates. The estimates obtained for the parameters γ and ρ are shown in Table 5, taken from [FLL15]. We note that the shape parameters ρ satisfy $-1 < \rho < 0$, which, as mentioned at the end of Section 2, implies that the Weibull survival functions can be expressed as mixtures of exponentials. In order to compare our approximation by exponential mixtures with the Weibull survival functions obtained in [FLL15], we used least-squares curve fitting to fit Weibulls (4) to the exponential mixtures obtained from (9) and (10), using the parameters from Table 2, in a similar manner to [JG10]. The estimated parameters γ and ρ are shown in Table 6, together with R^2 . It can be seen that these estimates are very close to those in Table 5.

Data set	γ	ρ	R^2
Business	1.5576	-0.8321	0.9818
Entertainment	1.7877	-0.7453	0.9957
Science	0.7575	-0.8318	0.9303
Shopping	0.4736	-0.7608	0.9240
Sports	3.1986	-0.8666	0.9967
Travel	3.5141	-0.9291	0.9760
All Categories	4.5262	-0.9134	0.9955

Table 5: Nonlinear least-squares regression to Weibull survival functions using the Kaplan-Meier estimates.

Finally, we note from Table 2 that λ_1 is significantly smaller than λ_2 . It therefore follows, from the comment made after (3) regarding the decreasing failure rate, that a substantial number of queries are popular for a short duration, because λ_2 is large, while the popularity of others lasts for much longer, because λ_1 is small. We note that a similar observation, in the context of the popularity of movies, was made in [PS10]. In particular, for the Google Trends data set, about 40% of the queries survive for only a single month.

5 Concluding Remarks

In this paper, we propose a stochastic evolutionary urn-based model that can be applied to survival analysis scenarios in the context of human dynamics. In our model, actors survive

Data set	γ	ρ	R^2
Business	1.5516	-0.8315	0.9967
Entertainment	1.7326	-0.7389	0.9976
Science	0.7500	-0.8276	0.9962
Shopping	0.4500	-0.7510	0.9930
Sports	3.0612	-0.8613	0.9971
Travel	3.2840	-0.9250	0.9969
All Categories	4.3860	-0.9108	0.9972

Table 6: Least-squares curve fitting of Weibulls to the exponential mixture survival functions.

in the system until they die and leave the system according to a constant mortality rate. A solution to the equations describing the model was obtained in (9), and it was seen that this solution is approximately an exponential mixture. We then successfully applied our model to data on the survival of popular search engine queries. It is also potentially applicable to other data sets relating to human behaviour.

One important problem to be addressed is how to choose the number of components in the mixture model. As discussed in [MP00, Chapter 6], several methods have been suggested to obtain a parsimonious model. Among these, a number of prominent methods are based on maximising a penalised log-likelihood function.

Generative models, like the one presented here, have the potential to explain observed social phenomena and, more specifically, social mechanisms and the emergence of collective behaviour, as discussed in the introduction. They allow us to gain insight into the underlying processes; in addition, they may be useful for more accurate prediction of, for example, online content, as discussed in [TdFA14]. With regards to our theme of human dynamics [CFL09], our stochastic mixture urn model contributes to a better understanding of collective phenomena, such as popularity, and how such global behaviour emerges from the decisions of individuals from whom the statistics have been collected. Another possible avenue for investigation is to consider other distribution mixtures, such as finite Weibull mixture models (see [BNF04], where it is argued that the reliability of a system can be suitably modelled by a Weibull mixture model).

Acknowledgements

We would like to thank Suneel Kingrani, who collected the Google Trends data and computed the Kaplan-Meier estimates for the data set.

References

- [AB02] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47–97, 2002.
- [Bar05] A.-L. Barabási. The origin of bursts and heavy tails in human dynamics. *Nature*, 435:207–211, 2005.

- [Bar07] A.-L. Barabási. The architecture of complexity: From network structure to human dynamics. *IEEE Control Systems Magazine*, 27:33–42, 2007.
- [BG03] C. Bracquemond and O. Gaudoin. A survey on discrete lifetime distributions. *International Journal of Reliability, Quality and Safety Engineering*, 10:69–98, 2003.
- [BJ97] H. Block and H. Joe. Tail behavior of the failure rate functions of mixtures. *Lifetime Data Analysis*, 3:269–288, 1997.
- [BNF04] T. Búcar, M. Nagode, and M. Fajdiga. Reliability approximation using finite Weibull mixture distributions. *Reliability Engineering and System Safety*, 84:241–251, 2004.
- [BSV07] S. Börner, S. Sanyal, and A. Vespignani. Network science. *Annual Review of Information Science & Technology (ARIST)*, 41:537–607, 2007.
- [CFL09] C. Castellano, S. Fortunato, and V. Loreto. Statistical physics of social dynamics. *Reviews of Modern Physics*, 81:591–646, 2009.
- [CGW⁺08] J. Candia, M.C. González, P. Wang, T. Schoenhar, G. Madey, and A.-L. Barabási. Uncovering individual and collective human dynamics from mobile phone records. *Journal of Physics A: Mathematical and Theoretical*, 41:224015, 11pp, 2008.
- [Fel71] W. Feller. *An Introduction to Probability Theory and its Applications*, volume II. John Wiley & Sons, New York, NY, 2nd edition, 1971.
- [Fin08] M. Finkelstein. *Failure Rate Modelling for Reliability and Risk*. Springer Series in Reliability Engineering. Springer-Verlag, London, 2008.
- [Fin09] M. Finkelstein. Understanding the shape of the mixture failure rate (with engineering and demographic applications). *Applied Stochastic Models In Business And Industry*, 25:643–663, 2009.
- [FL00] T.R. Fleming and D.Y. Lin. Survival analysis in clinical trials: Past developments and future directions. *Biometrics*, 56:971–983, 2000.
- [FLL07] T. Fenner, M. Levene, and G. Loizou. A model for collaboration networks giving rise to a power-law distribution with an exponential cutoff. *Social Networks*, 29:70–80, 2007.
- [FLL12] T. Fenner, M. Levene, and G. Loizou. A discrete evolutionary model for chess players ratings. *IEEE Transactions on Computational Intelligence and AI in Games*, 4:84–93, 2012.
- [FLL13] T. Fenner, M. Levene, and G. Loizou. A bi-logistic growth model for conference registration with an early bird deadline. *Central European Journal of Physics*, 11:904–909, 2013.
- [FLL15] T. Fenner, M. Levene, and G. Loizou. A stochastic evolutionary model for capturing human dynamics. *Journal of Statistical Mechanics: Theory and Experiment*, August:P08015, 2015.

- [FW98] A. Feldman and W. Whitt. Fitting mixtures of exponentials to long-tail distributions to analyze network performance models. *Performance Evaluation*, 31:245–279, 1998.
- [Gal08] S. Galam. Sociophysics: A review of Galam models. *Journal of Modern Physics C*, 19:409–440, 2008.
- [GG01] L.A. Gavrilov and N.S. Gavrilova. The reliability theory of aging and longevity. *Journal of Theoretical Biology*, 213:527–545, 2001.
- [HS98] P. Hedström and R. Swedberg. Social mechanisms: An introductory essay. In P. Hedström and R. Swedberg, editors, *Social Mechanisms: An Analytical Approach to Social Theory*, pages 1–31. Cambridge University Press, Cambridge, U.K., 1998.
- [Jew82] N.P. Jewell. Mixture of exponential distributions. *Annals of Statistics*, 10:479–484, 1982.
- [JG10] T. Jin and L.K. Gonigunta. Exponential approximation of Weibull renewal with decreasing failure rate. *Journal of Statistical Computation and Simulation*, 80:273–285, 2010.
- [Kea99] C. Keatinge. Modeling losses with the mixed exponential distribution. *Proceedings of the Casualty Actuarial Society*, LXXXVI:654–698, 1999.
- [KK12] D.G. Kleinbaum and M. Klein. *Survival Analysis, A Self-Learning Text*. Springer Science+Business Media, LLC, New York, NY, 3rd edition, 2012.
- [KM58] E.L. Kaplan and P. Meier. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53:457–481, 1958.
- [LAE⁺12] E. Lovecchio, P. Allegrini, E. Geneston, B.J. West, and P. Grigolini. From self-organized to extended criticality. *Frontiers in Physiology*, 3:Article 98, 9pp, 2012.
- [Law03] J.F. Lawless. *Statistical Models and Methods for Lifetime Data*. Wiley Series in Probability and Statistics. John Wiley & Sons, 2nd edition, 2003.
- [LMS12] J.G. Lee, S. Moon, and K. Salamatian. Modeling and predicting the popularity of online contents with Cox proportional hazard regression model. *Neurocomputing*, 76:134–145, 2012.
- [MAAJ13] J. Mathiesen, L. Angheluta, P.T.H. Ahlgren, and M.H. Jensen. Excitable human dynamics driven by extrinsic events in massive communities. *Proceedings of the National Academy of Sciences of the United States of America*, 110:17259–17262, October 2013.
- [MD09] A.T. Mendoza-Rosas and S. De la Cruz-Reyna. A mixture of exponentials distribution for a simple and precise assessment of the volcanic hazard. *Natural Hazards and Earth System Sciences*, 9:425–431, 2009.
- [Mot95] H. Motulsky. *Intuitive Biostatistics*. Oxford University Press, Oxford, 1995.

- [MP00] G. McLachlan and D. Peel. *Finite Mixture Models*. John Wiley & Sons, New York, NY, 2nd edition, 2000.
- [MSR13] A. Mitnitski, X. Song, and K. Rockwood. Assessing biological aging: the origin of deficit accumulation. *Biogerontology*, 14:709–717, 2013.
- [OK12] P.D.T. O’Connor and A. Kleyner. *Practical Reliability Engineering*. Wiley Series in Telecommunications. John Wiley & Sons, Chichester, 5th edition, 2012.
- [Pro63] F. Proschan. Theoretical explanation of observed decreasing failure rate. *Technometrics*, 5:375–383, 1963.
- [PS10] R.K. Pan and S. Sinha. The statistical laws of popularity: universal properties of the box-office dynamics of motion pictures. *New Journal of Physics*, 12:115004 (23pp), 2010.
- [Rin09] H. Rinne. *The Weibull Distribution: A Handbook*. CRC Press, Boca Raton, FL, 2009.
- [RNP⁺10] J.T. Rich, J.G. Neely, R.C. Paniello, C.C.J. Voelker, B. Nussenbaum, and E.W. Wang. A practical guide to understanding Kaplan-Meier curves. *Otolaryngology-Head and Neck Surgery*, 143:331–336, 2010.
- [SBCG13] C.M. Schneider, V. Belik, T. Couronné, and M.C. González. Unravelling daily human mobility motifs. *Journal of the Royal Society Interface*, 10:20130246, 2013.
- [SC88] S. Siegel and N.J. Castellan Jr. *Nonparametric Statistics for the Behavioral Sciences*. McGraw-Hill, New York, NY, 2nd edition, 1988.
- [SC14] P. Sen and B.K. Chakrabarti. *Sociophysics: An Introduction*. Oxford University Press, Oxford, 2014.
- [Sch98] T.C. Schelling. Social mechanisms and social dynamics. In P. Hedström and R. Swedberg, editors, *Social Mechanisms: An Analytical Approach to Social Theory*, pages 32–44. Cambridge University Press, Cambridge, U.K., 1998.
- [Sim55] H.A. Simon. On a class of skew distribution functions. *Biometrika*, 42:425–440, 1955.
- [TdFA14] A. Tatar, M.D. de Amorim, S. Fdida, and P. Antoniadis. A survey on predicting the popularity of web content. *Journal of Internet Services and Applications*, 5:20 pages, 2014.
- [TW02] A. Tsoularis and J. Wallace. Analysis of logistic growth models. *Mathematical Biosciences*, 179:21–55, 2002.
- [VAFL10] P.O.S. Vaz De Melo, L. Akoglu, C. Faloutsos, and A.A.P. Loureiro. Surprising patterns for the call duration distribution of mobile phone users. In *Proceedings of European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)*, pages 354–369, Barcelona, 2010.

- [ZPL12] D. Zhang, K. Prior, and M. Levene. How long do wikipedia editors keep active? In *Proceedings of the 8th International Symposium on Wikis and Open Collaboration (WikiSym)*, Linz, Austria, 2012.