

Burstiness Scale: A Parsimonious Model for Characterizing Random Series of Events

Rodrigo A S Alves
Department of Applied
Social Sciences
CEFET-MG
rodrigo@dcsa.cefetmg.br

Renato Assunção
Department of
Computer Science
UFMG
assuncao@dcc.ufmg.br

Pedro O.S. Vaz de Melo
Department of
Computer Science
UFMG
olmo@dcc.ufmg.br

ABSTRACT

The problem to accurately and parsimoniously characterize random series of events (RSEs) seen in the Web, such as Yelp reviews or Twitter hashtags, is not trivial. Reports found in the literature reveal two apparent conflicting visions of how RSEs should be modeled. From one side, the Poissonian processes, of which consecutive events follow each other at a relatively regular time and should not be correlated. On the other side, the self-exciting processes, which are able to generate bursts of correlated events. The existence of many and sometimes conflicting approaches to model RSEs is a consequence of the unpredictability of the aggregated dynamics of our individual and routine activities, which sometimes show simple patterns, but sometimes results in irregular rising and falling trends. In this paper we propose a parsimonious way to characterize general RSEs, namely the *Burstiness Scale (BuSca)* model. *BuSca* views each RSE as a mix of two independent processes: a *Poissonian* and a self-exciting one. Here we describe a fast method to extract the two parameters of *BuSca* that, together, gives the burstiness scale ψ , which represents how much of the RSE is due to bursty and viral effects. We validated our method in eight diverse and large datasets containing real random series of events seen in Twitter, Yelp, e-mail conversations, Digg, and online forums. Results showed that, even using only two parameters, *BuSca* is able to accurately describe RSEs seen in these diverse systems, what can leverage many applications.

Keywords

communication dynamics; self-exciting point process; social media

1. INTRODUCTION

What is the best way to characterize random series of events (RSEs) seen on the Web, such as Yelp reviews or Twitter hashtags? Descriptively, one can characterize a given RSE as constant for a period, then bursty for another, back to being constant and, after a long period, bursty again. Formally, the answer to this question is not trivial. It certainly must include the extreme case of the homogeneous *Poisson Process* (PP) [6], which has a single and intuitive rate parameter λ . Consecutive events of PP follow each other at a

relatively regular time and λ represents the constant rate at which events arrive. The class of *Poissonian* or *completely random* processes includes also the case when λ varies with time. In this class, events must be without any aftereffects, i.e., there is no interaction between any sequence of events [17]. There are RSEs seen on the Web that were accurately modeled by a Poissonian process, such as many instances of viewing activity on Youtube [3], e-mail conversations [12] and hashtag posts on Twitter [11].

Unfortunately, recent analyses showed that this simple and elegant model has proved unsuitable for many cases [14, 21, 19, 20]. Such analyses revealed that many RSEs produced by humans have very long periods of inactivity and also bursts of intense activity [1, 7], in contrast to Poissonian processes, where activities may occur at a fairly constant rate. Moreover, many RSEs on the Web also have strong correlations between historical data and future data [3, 19, 20, 4], a feature that must not occur in Poissonian processes. These RSEs fall into a particular class of random point processes, the so called *self-exciting* processes [17]. The problem of characterizing such RSEs is that they occur in many shapes and in very unpredictable ways [20, 13, 5, 4, 24, 3]. They have the so called “quick rise-and-fall” property [13] of bursts in cascades, producing correlations between past and future data that are not captured by regression-based methods [23].

As pointed out by [3], the aggregated dynamics of our individual activities is a consequence of a myriad of factors that guide individual actions, producing a plethora of collective behaviors. In order to accurately capture all patterns seen in human-generated RSEs, researchers are proposing models with many parameters and, for most of the times, tailored to a specific activity in a specific system [12, 24, 11, 26, 23, 13, 5, 4, 16, 27, 25]. Going against this trend, in this work we propose the *Burstiness Scale (BuSca)* model, a parsimonious model to characterize RSEs that can be (i) purely Poissonian, (ii) purely self-exciting or (iii) a mix of these two behaviors. In *BuSca*, the underlying Poissonian process is responsible for the arrival of events related to the routine activity dynamics of individuals, whereas the underlying self-exciting process is responsible for the arrival of bursty and ephemeral events, related to the endogenous (e.g. online social networks) and the exogenous (e.g. mass media) mechanisms that drive public attention and generate the “quick rise-and-fall” property and correlations seen in many RSEs [3, 11]. To illustrate that, observe Figure 1, which shows the cumulative number of occurrences $N(t)$ of three Twitter hashtags over time. In Figure 1a, the curve of hashtag #wheretheydothatat is a straight line, indicating that this RSE is well modeled by a PP. In Figure 1b, the curve of hashtag #cotto has long periods of inactivities and a burst of events, suggesting that the underlying process may be self-exciting in this

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '16, August 13-17, 2016, San Francisco, CA, USA

© 2016 ACM. ISBN 978-1-4503-4232-2/16/08...\$15.00

DOI: <http://dx.doi.org/10.1145/2939672.2939852>

case. Finally, In Figure 1c, the curve of hashtag #t_a is apparently a mix of these two processes, exactly what *BuSca* aims to model.

Besides that, our goal is also to characterize general RSEs using the least amount of parameters possible. The idea is to propose a parsimonious model that can separate out constant and routine events from bursty and ephemeral events in general RSEs. We present and validate a particular and parsimonious case of *BuSca*, where the Poissonian process is given by a homogeneous Poisson process and the self-exciting process is given by a Self-Feeding Process (SFP) [20]. We chose these models because (i) both of them require a single parameter and (ii) they are on opposite ends of the spectrum. The PP is on the extreme side where the events do not interact with each other and inter-event times are independent. On the other extreme lies the SFP, where consecutive inter-event times are highly correlated. Even though *BuSca* has only two parameters, we show that, surprisingly, it is able to accurately characterize a large corpus of diverse RSE seen in Web systems, namely Twitter, Yelp, e-mails, Digg, and online forums. We show that disentangling constant from ephemeral events in general RSEs may reveal interesting, relevant and fascinating properties about the underlying dynamics of the system in question in a very summarized way, leveraging applications such as monitoring systems, anomaly detection methods, flow predictors, among others. In summary, the main contributions of this paper are:

- *BuSca*, a widely applicable model that characterize communication time series with only two intuitive parameters and validated in eight different datasets. We calculate the *burstiness scale* ψ , which represents how much of the process is due to bursty events.
- A fast and scalable method to separate events arising from a homogeneous Poisson Process from those arising from a self-exciting process in RSEs.
- A method to detect anomalies and another to detect bursts in random series of events.

The rest of the paper is organized as follows. In Section 2, we provide a brief survey of the related work. Our model is introduced in Section 3 together with the algorithm to estimate its parameters. We show that the maximum likelihood estimator is biased and show to fix the problem in Section 4, discussing a statistical test procedure to discriminate between extreme cases in Section 5. In Section 6, we describe the eight datasets used in this work and show the goodness of fit of our model in Section 7. A comparison with the Hawkes model is given in Section 8. We show two applications of our model in Section 9. We close the paper with Section 10, where we present our conclusions.

2. RELATED WORK

Characterizing the dynamics of human activity in the Web has attracted the attention of the research community as it has implications that can benefit a large number of applications, such as trend detection, popularity prediction, clustering, anomaly detection, among others [1, 3, 12, 22, 24, 11, 26, 23, 13, 5, 4, 16, 27, 25]. The problem is that uncovering the rules that govern human behavior is a difficult task, since many factors may influence an individual's decision to take action. Analysis of real data have shown that human activity in the Web can be highly unpredictable, ranging from being completely random [3, 2, 8, 9, 12] to highly correlated and bursty [1, 7, 20, 13, 25, 16, 27].

As one of the first attempts to model bursty RSEs, Barabási et al. [1] proposed that bursts and heavy-tails in human activities are

a consequence of a decision-based queuing process, when tasks are executed according to some perceived priority. The queuing models generate power law distributions, but do not correlate the timing of events explicitly. As an alternative to queuing models, many researchers started to consider the self-exciting point processes, which are also able to model correlations between historical and future events. In a pioneer effort, Crane and Sornette [3] modeled the viewing activity on Youtube as a Hawkes process. They proposed that the burstiness seen in data is a response to endogenous word-of-mouth effects or sudden exogenous perturbations. This seminal paper inspired many other efforts to model human dynamics in the Web as a Hawkes process [13, 25, 16, 27]. Similar to the Hawkes process, the Self-Feeding process (SFP) [20] is another type of self-exciting process that also captures correlations between historical and future data, being also used to model human dynamics in the Web [20, 19, 4]. Different from Hawkes, whose conditional intensity explicitly depends on all previous events, the SFP considers only the previous inter-event time.

Although there are strong evidences that self-exciting processes are well suited to model human dynamics in the Web, there are studies that show that the Poisson process and its variations are also appropriate [3, 2, 8, 9, 12]. [2, 8] showed that Internet traffic can be accurately modeled by a Poisson process under particular circumstances, e.g. heavy traffic. When analyzing Youtube viewing activity, [3] verified that 90% of the videos analyzed either do not experience much activity or can be described statistically as a Poisson process. Malmgreen et al. [12] showed that a non-homogeneous Poisson process can accurately describe e-mail communications. In this case, the rate $\lambda(t)$ varies with time, in a periodic fashion (e.g., people answer e-mails in the morning; then go to lunch; then answer more e-mails, etc).

These apparently conflicting approaches, i.e., self-exciting and Poissonian approaches, motivated many researchers to investigate and characterize this plethora of human behaviors. For instance, [18] used a machine learning approach to characterize videos on Youtube. From several features extracted from Youtube and Twitter, the authors verified that the current tweeting rate along with the volume of tweets since the video was uploaded are the two most important Twitter features for classifying a Youtube video into viral or popular. In this direction, [11] verified that Twitter hashtag activities may be continuous, periodic or concentrated around an isolated peak, while [5] found that revisits account from 40% to 96% of the popularity of an object in Youtube, Twitter and LastFm, depending on the application. [26] verified that the popularity of a Youtube video can go through multiple phases of rise and fall, probably generated by a number of different background random processes that are super-imposed onto the power-law behavior. The main difference between these models and ours is that the former ones mainly focus on representing all the details and random aspects of very distinct RSEs, which naturally demands many parameters. In our case, our proposal aims to disentangle the bursty and constant behavior of RSEs as parsimoniously as possible. Surprisingly, our model is able to accurately describe a large and diverse corpus of RSEs seen in the Web with only two parameters.

In this version, random series of events are modeled by a mixture of two independent processes: a Poisson process, which accounts for the background constant behavior, and a one-parameter SFP, which accounts for the bursty behavior. A natural question that arises is: how different is this model from the widely used Hawkes process? The main difference are twofold. First, in the Hawkes process, every single arriving event excites the process, i.e., is correlated to the appearance of future events. In our proposal, since the PP is independent, non-correlated events may arrive at any time.

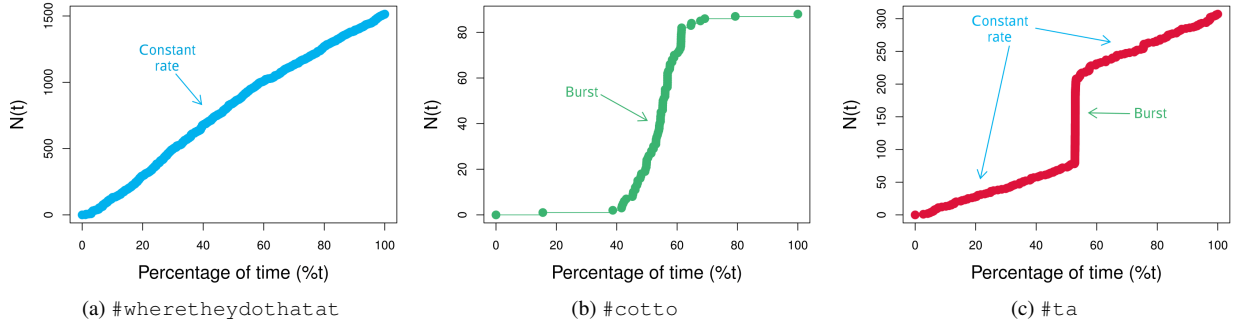


Figure 1: Three real individuals from twitter database.

Second, our model is even more parsimonious than the Hawkes process, two parameters against three¹. In Section 8 we quantitatively show that our proposed model is more suited to real data than the Hawkes process.

3. MODELING INFORMATION BURSTS

Point processes is the stochastic process framework developed to model a random sequence of events (RSE). Let $0 < t_1 < t_2 < \dots$ be a sequence of random event times, with $t_i \in \mathbb{R}^+$, and $N(a, b)$ be the random number of events in $(a, b]$. We simplify the notation when the interval starts at $t = 0$ by writing simply $N(0, b) = N(b)$. Let \mathcal{H}_t be the random history of the process up to, but not including, time t . A fundamental tool for modeling and for inference in point processes is the conditional intensity rate function. It completely characterizes the distribution of the point process and it is given by

$$\lambda(t|\mathcal{H}_t) = \lim_{h \rightarrow 0} \frac{\mathbb{E}(N(t, t+h)|\mathcal{H}_t)}{h}. \quad (1)$$

The interpretation of this random function is that, for a small time interval h , the value of $\lambda(t|\mathcal{H}_t) \times h$ is approximately the expected number of events in $(t, t+h)$. The notation emphasizes that the conditional intensity at time t depends on the random events that occur previous to t .

The most well known point process is the Poisson process where $\lambda(t|\mathcal{H}_t) = \lambda(t)$, a deterministic function. The main characteristic of a Poisson process is that the counts in disjoint intervals are independent random variables. When the intensity does not vary in time, with $\lambda(t) \equiv \lambda$, we have a homogeneous Poisson process.

3.1 Self-feeding process

The self-feeding process (SFP) [20, 19] conditional intensity has a simple dependence on its past. Locally, it acts as a homogeneous Poisson process but its conditional intensity rate is inversely proportional to the temporal gap between the two last events. More specifically, the conditional intensity function is given by

$$\lambda_s(t|\mathcal{H}_t) = \frac{1}{\mu/e + \Delta t_i} \quad (2)$$

where $\Delta t_i = t_i - t_{i-1}$ and $t_i = \max_k\{t_k : t_k \leq t\}$. This implies that the inter-event times $\Delta t_{i+1} = t_{i+1} - t_i$ are exponentially distributed with expected value $\mu/e + \Delta t_i$. The inter-event times Δt_i follow a Markovian property. The constant μ is the median of the inter-event times and $e \approx 2.718$ is the Euler constant. A more general version of the SFP uses an additional parameter ρ which was taken equal to 1 in this work. The motivation for this is

¹Considering the most parsimonious version of the Hawkes process.

that, in many databases analysed previously [20, 19], it was found that $\rho \approx 1$. An additional benefit of this decision is the simpler likelihood calculations involving the SFP.

The Figure 2a presents three realizations of the SFP process in the interval $(0, 100)$ with parameter $\mu = 1$. The vertical axis shows the accumulated number of events $N(t)$ up to time t . One striking aspect of this plot is its variability. In the first 40 time units, the lightest individual shows a rate of approximately 0.5 events per unit time while the darkest one has a rate of 2.25. Having accumulated a very different number of events, they do not have many additional points after time $t = 40$. The third one has a more constant rate of increase during the whole time period. Hence, with the same parameter μ , we can see very different realizations from the SFP process. A common characteristic of the SFP instances is the mix of bursty periods alternating with quiet intervals.




Figure 2: Three instances of the SFP process with $\mu = 1$ in the time interval $[0, 100)$ (left) and a real time series with the hashtag `shaq` from Twitter.

However, in many datasets we have observed a departure from this SFP behavior. The most noticeable discrepancy is the absence of the long quiet periods predicted by the SFP model. To be concrete, consider the point processes realizations in Figure 2b. This plot is the cumulative counting of Twitter posts from the hashtag `shaq`. There are two clear bursts, when the series has a large increase in a short period of time. Apart from these two periods, the counts increase in a regular and almost constant rate. We do not observe long stretches of time with no events, as one would expect to see if a SFP process is generating these data.

3.2 The BuSca Model

We propose a point process model that exhibits the same behavior consistently observed in our empirical findings: we want a mix of random bursts followed by more quiet periods, and we want realizations where the long silent periods predicted by the SFP are not allowed. To obtain these two aspects we propose a new model that is a mixture of the SFP process, to guarantee the presence of random bursts, with a homogeneous Poisson process, to generate a random but rather constant rate of events, breaking the long empty spaces created by the SFP. While the SFP captures the viral and

ephemeral “rise and fall” patters, the PP captures the routine activities, acting as a random background noise added to a signal point process. We call this model the *Burstiness Scale (BuSca)* model.

Figure 3 shows the main idea of *BuSca*. The observed events are those on the bottom line. They are composed by two types of events, each one generated by a different point process. Each observed event can come either from a Poisson process (top line) or from an SFP process (middle line). We observe the mixture of these two types of events on the third line without an identifying label. This lack of knowledge of the source process for each event is the cause of most inferential difficulties, as we discuss later.




Figure 3: The *BuSca* model. The top line displays the events from the Poisson(λ_p) component along the timeline while the middle line displays those from the SFP(μ) component. The user observes only the third line, the combination of the first two, without a label to identify the source process associated with each event.

Figure 4 shows different realizations of the mixture process in the time interval $[0, 100]$. In each plot, the curves show the cumulative number of events up to time t . The blue line represents a homogeneous Poisson process realization with parameter λ_p while the green curve represents the SFP with parameter μ . The red curve represents the mixture of the two other realizations.

3.3 The likelihood function

The log-likelihood function $\ell(\theta)$ for any point process is a function of the conditional intensity $\lambda(t|\mathcal{H}_t)$ and of the events $t_1 < t_2 < \dots$:

$$\ell(\theta) = \left(\sum_{i=1}^n \log \lambda(t_i|\mathcal{H}_{t_i}) \right) - \int_a^b \lambda(t|\mathcal{H}_t) dt \quad (3)$$

The conditional intensity function $\lambda_m(t|\mathcal{H}_t)$ of *BuSca* is the sum of the conditional intensities of the component processes, the Poisson intensity $\lambda_p(t|\mathcal{H}_t) = \lambda_p$, and the SFP intensity $\lambda_s(t|\mathcal{H}_t)$:

$$\lambda_m(t|\mathcal{H}_t) = \lambda_p + \lambda_s(t|\mathcal{H}_t) \quad (4)$$

The stochastic history \mathcal{H}_t of the mixed process contains only the events’ times t_1, t_2, \dots but not their identifying component processes labels, either s (from SFP) or p (from PP), for each event. The log-likelihood function for the mixture process observed in the time interval $[a, b]$ is given by

$$\ell(\theta) = \sum_{i=1}^n \log [\lambda_p + \lambda_s(t_i|\mathcal{H}_{t_i})] - \int_a^b \lambda_s(t|\mathcal{H}_t) dt - (b-a)\lambda_p \quad (5)$$

The log-likelihood (5) is not computable because $\lambda_s(t|\mathcal{H}_t)$ requires the knowledge of the last SFP inter-event time Δt_i for each $t \in (a, b]$. This would be known only if the source-process label for each event in the observed mixture is also known. Since these labels are hidden, we adopt the EM algorithm to obtain the maximum likelihood estimates $\hat{\lambda}_p$ and $\hat{\mu}$. We define the *burstiness scale* $\psi = [(1 - \lambda_p/(\lambda_p + (b-a)/\mu)) 100\%]$ as the percentage of bursty events in a given RSE. It can be estimated by $\hat{\psi} = [(1 - \hat{\lambda}_p(b-a)/n) 100\%]$. This gives the estimated proportion of events that comes from the pure SFP process. The latent labels are also inferred as part of the inferential procedure.

The use of the EM algorithm in the case of point process mixtures is new and presents several special challenges with respect to the usual EM method. The reason for the difficulty is the lack of independent pieces in the likelihood. The correlated sequential data in the likelihood brings several complications dealt with in the next two sections. The final algorithm has a complexity of $O(n^2)$, as can be seen in the extended version of this paper.

3.4 The E step

The EM algorithm requires the calculation of $\mathbb{E}[\ell(\theta)]$, the expected value of the log-likelihood (5) with respect to the hidden labels. Using a Taylor expansion,

$$\mathbb{E}[\log(X)] \approx \log \mathbb{E}[X] - \frac{\mathbb{V}[X]}{2\mathbb{E}[X]^2} = \log \mathbb{E}[X] - \frac{\mathbb{E}[X^2] - \mathbb{E}[X]^2}{2\mathbb{E}[X]^2}$$

we have

$$\begin{aligned} \mathbb{E}[\ell(\theta)] \approx & \sum_{i=1}^n \left[\log(\lambda_p + \mathbb{E}[\lambda_s(t_i|\mathcal{H}_{t_i})]) \right. \\ & \left. - \frac{\mathbb{E}[\lambda_s(t_i|\mathcal{H}_{t_i})^2] - \mathbb{E}[\lambda_s(t_i|\mathcal{H}_{t_i})]^2}{2(\lambda_p + \mathbb{E}[\lambda_s(t_i|\mathcal{H}_{t_i})])^2} \right] - \int_a^b \mathbb{E}[\lambda_s(t|\mathcal{H}_t)] dt - (b-a)\lambda_p \end{aligned} \quad (6)$$

A naive way to obtain the required

$$\mathbb{E}[\lambda_s(t|\mathcal{H}_t)] = \mathbb{E} \left[\frac{1}{\mu/e + \Delta t_i} \right] \quad (7)$$

is to consider all possible label assignments to the events and its associated probabilities. Knowing which events belong to the SFP component, we also know the value of Δt_i . Hence, it is trivial to evaluate $1/(\mu/e + \Delta t_i)$ in each one these assignments for any t , and finally to obtain (7) by summing up all these values multiplied by the corresponding label assignment probabilities. This is unfeasible because the number of label assignments is too large, unless the number of events is unrealistically small.

To overcome this difficulty, we developed a dynamic programming algorithm. Figure (5) shows the conditional intensity $\lambda_s(t|\mathcal{H}_t)$ up to, and not including, t_{i-1} as green line segments and the constant Poisson intensity λ_p as a blue line. Our algorithm is based on a fundamental observation: if t_{i-1} comes from the Poisson process, it does not change the current SFP conditional intensity until the next event t_i comes in.

We start by calculating $a_i \equiv \mathbb{E}[\lambda_s(t_i|\mathcal{H}_{t_i})]$ conditioning on the t_{i-1} event label:

$$\begin{aligned} a_i = \mathbb{E}[\lambda_s(t_i|\mathcal{H}_{t_i})] = & \mathbb{P}(t_{i-1} \in \text{Poisson}|\mathcal{H}_{t_i}) \mathbb{E}[\lambda_s(t_i|\mathcal{H}_{t_i}, t_{i-1} \in \text{Poisson})] \\ & + \mathbb{P}(t_{i-1} \in \text{SFP}|\mathcal{H}_{t_i}) \mathbb{E}[\lambda_s(t_i|\mathcal{H}_{t_i}, t_{i-1} \in \text{SFP})] \end{aligned} \quad (8)$$

The evaluation of $\lambda_s(t_i|\mathcal{H}_{t_i})$ depends on the label assigned to t_{i-1} . If $t_{i-1} \in \text{Poisson}$, as in Figure 6, the last SFP inter-event time interval is the same as that for $t_{i-2} \leq t < t_{i-1}$, since a Poisson event does not change the SFP conditional intensity. Therefore, in this case,

$$\mathbb{E}[\lambda_s(t_i|\mathcal{H}_{t_i}, t_{i-1} \in \text{Poisson})] = \mathbb{E}[\lambda_s(t_{i-1}|\mathcal{H}_{t_{i-1}})] = a_{i-1}. \quad (9)$$

For the integral component in (6), we need the conditional intensity for t in the continuous interval and not only at the observed t_i values. However, by the same argument used for t_i , we have

$$\lambda_s(t|\mathcal{H}_t, t_{i-1} \in \text{Poisson}) = \lambda_s(t_{i-1}|\mathcal{H}_{t_{i-1}})$$

for $t \in [t_{i-1}, t_i)$.

The probability that the $(i-1)$ -th mixture event comes from the SFP component is proportional to its conditional intensity at the

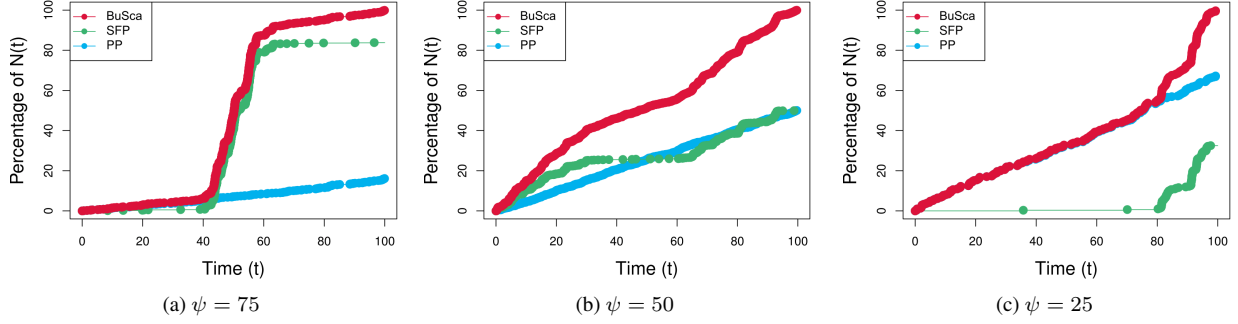


Figure 4: Realizations of the mixture process with different values for ψ in the time interval $[0,100]$.

event time t_i :

$$\mathbb{P}(t_{i-1} \in \text{SFP} | \mathcal{H}_{t_i}) = \frac{\lambda_s(t_{i-1} | \mathcal{H}_{t_i})}{\lambda_s(t_{i-1} | \mathcal{H}_{t_i}) + \lambda_p} \approx \frac{a_{i-1}}{a_{i-1} + \lambda_p}. \quad (10)$$

Therefore, using (9) and (10), we can rewrite (8) approximately as a recurrence relationship:

$$a_i = \frac{\lambda_p}{a_{i-1} + \lambda_p} a_{i-1} + \frac{a_{i-1}}{a_{i-1} + \lambda_p} \mathbb{E}[\lambda_s(t_i | \mathcal{H}_{t_i}, t_{i-1} \in \text{SFP})]. \quad (11)$$

We turn now to explain how to obtain the last term in (11). If $t_{i-1} \in \text{SFP}$, as exemplified in Figure 7, the last SFP inter-event time must be updated and it will depend on the most recent SFP event previous to t_{i-1} . There are only $i-2$ possibilities for this last previous SFP event and this fact is explored in our dynamic programming algorithm. Recursively, we condition on these possible $i-2$ possibilities to evaluate the last term in (11). More specifically, the value of $\mathbb{E}[\lambda_s(t_i | \mathcal{H}_{t_i}, t_{i-1} \in \text{SFP})]$ is given by

$$\begin{aligned} & \mathbb{E}[\lambda_s(t_i | \mathcal{H}_{t_i}, (t_{i-1}, t_{i-2}) \in \text{SFP}) \mathbb{P}(t_{i-2} \in \text{SFP} | \mathcal{H}_{t_i}, t_{i-1} \in \text{SFP}) \\ & + \mathbb{E}[\lambda_s(t_i | \mathcal{H}_{t_i}, t_{i-1} \in \text{SFP}, t_{i-2} \in \text{Poisson})] \\ & \mathbb{P}(t_{i-2} \in \text{Poisson} | \mathcal{H}_{t_i}, t_{i-1} \in \text{SFP}) \\ & \approx \frac{1}{\mu/e + (t_{i-1} - t_{i-2})} \frac{a_{i-2}}{a_{i-2} + \lambda_p} \\ & + \mathbb{E}[\lambda_s(t_i | \mathcal{H}_{t_i}, t_{i-1} \in \text{SFP}, t_{i-2} \in \text{Poisson})] \frac{\lambda_p}{a_{i-2} + \lambda_p} \end{aligned} \quad (12)$$

When the last two events t_{i-1} and t_{i-2} come from the SFP process, we know that the conditional intensity of the SFP process is given by the first term in (12). The unknown expectation in (12) is obtained by conditioning in the t_{i-3} label. In this way, we recursively walk backwards, always depending on one single unknown of the form

$$\mathbb{E}[\lambda_s(t_i | \mathcal{H}_{t_i}, t_{i-1} \in \text{SFP}, \{t_{i-2}, t_{i-3}, \dots, t_k\} \in \text{Poisson})]$$

where $k < i-2$. At last, we can calculate a_i in (11) by the iterative expression

$$\mathbb{E}[\lambda_s(t_i | \mathcal{H}_{t_i}, t_{i-1} \in \text{SFP})] = \frac{a_{i-1}}{a_{i-1} + \lambda_p} \sum_{k=1}^{i-2} \left\{ \frac{a_k}{a_k + \lambda_p} \frac{1}{\mu/e + (t_{i-1} - t_k)} \prod_{j=k+1}^{i-2} \frac{\lambda_p}{a_j + \lambda_p} \right\} \quad (13)$$

We have more than one option as initial conditions for this iterative computation. One is to assume that the first two events belong to the SFP. Another one is to use

$$\lambda_s(t_i | \mathcal{H}_{t_i}, t_{i-1} \in \text{SFP}, \{t_{i-2}, t_{i-3}, \dots, t_2\} \in \text{Poisson}) = \frac{1}{\mu/e + \mu}$$

and the first event comes from the SFP. Even with a moderate num-




Figure 5: Start




Figure 6: $t_{i-1} \in \text{Poisson}$

ber of events, this initial condition choices affect very little the final results and either of them can be selected in any case.

To end the E-step, the log-likelihood in (6) requires also $\mathbb{E}[\lambda_s(t_i | \mathcal{H}_{t_i})^2]$. This is calculated in an entirely analogous way as we did above.

3.5 The M step

Different from the E step, the M step did not require special development from us. Having obtained the log-likelihood (6) we simply maximize the likelihood and update the estimated parameter values of $\hat{\mu}$ and $\hat{\lambda}_p$. In this maximization procedure, we constrain the search within two intervals. For an observed point pattern with n events, we use $[0, n/t_n]$ for λ_p . The intensity must be positive and hence, the zero lower bound represents a pure SFP process while the the upper bound represents the maximum likelihood estimate of λ_p in the other extreme case of a pure homogeneous Poisson process. For the μ parameter, we adopt the search interval $[0, t_n]$. Since μ is the median inter-event time in the SFP component, a value $\mu \approx 0$ induces a pattern with a very large of events while $\mu = t_n$ represents, in practice, a pattern with no SFP events.

4. MLE BIAS AND A REMEDY

Several simulations were performed to verify that the estimation of the parameters proposed in Section 3 is suitable. There is no theory about the MLE behavior in the case of point processes data following a complex mixture model as ours. For this, synthetic data were generated by varying the sample size n and the parameters λ_p and μ of the mixture. We vary n in $\{100, 200, \dots, 1000\}$ for each




Figure 7: $t_{i-1} \in \text{SFP}$

pair (λ_p, μ) . The parameters λ_p and μ were empirically selected in such a way that the expected percentage of points coming from the SFP process (denoted by the burstiness scale parameter ψ) varied in $\{10\%, 20\%, \dots, 90\%\}$.

For each pair (n, ψ) , we conducted 100 simulations, totaling 9,000 simulations. In each simulation, the estimated parameters $(\hat{\lambda}_p, \hat{\mu})$ were calculated by the EM algorithm. Since their range vary along the simulations, we considered their relative differences with respect to the true values (λ_p, μ) . For μ , define

$$\Delta(\hat{\mu}, \mu) = \begin{cases} \hat{\mu}/\mu - 1, & \text{if } \hat{\mu}/\mu \geq 1 \\ 1 - \mu/\hat{\mu}, & \text{otherwise} \end{cases} \quad (14)$$

The main objective of this measure is to treat symmetrically the relative differences between $\hat{\mu}$ and μ . Consider a situation where $\Delta(\hat{\mu}, \mu) = 0.5$. This means that $\hat{\mu} = 1.5\mu$. Symmetrically, if $\mu = 1.5\hat{\mu}$, we have $\Delta(\hat{\mu}, \mu) = -0.5$. The value $\Delta(\hat{\mu}, \mu) = 0$ implies that $\hat{\mu} = \mu$. We define $\Delta(\hat{\lambda}_p, \lambda_p)$ analogously.

4.1 The estimator $\hat{\lambda}_p$

The results of $\Delta(\hat{\lambda}_p, \lambda_p)$ are shown in Figure 8. Each boxplot correspond to one of the 90 possible combinations of (ψ, n) . The vertical blue lines separate out the different values of ψ . Hence, the first 10 boxplots are those calculated to the combinations $(\psi = 90\%, n = 100, 200, \dots, 1000)$. The next 10 boxplots correspond to the values of $\Delta(\hat{\lambda}_p, \lambda_p)$ for $\psi = 80\%$ and $n = 100, 200, \dots, 1000$. The absolute values for $\Delta(\hat{\lambda}_p, \lambda_p)$ were censored at 5 and this is represented by the horizontal red lines at heights -5 and $+5$.




Figure 8: Boxplots of $\Delta(\hat{\lambda}_p, \lambda_p)$ according to $(100 - \psi)$ and n .

The estimator $\hat{\lambda}_p$ is well behaved, with small bias and variance decreasing with the sample size. The only cases where it has a large variance is when the total sample size is very small and, at the same time, the percentage of Poisson events is also very small. For example, with $n = 200$ and $\psi = 90\%$, we expect to have only 20 *unidentified* Poisson cases and it is not reasonable to expect an accurate estimate in this situation.

4.2 The estimator $\hat{\mu}$

The estimator $\hat{\mu}$ overestimates the true value of μ in all cases with the bias increasing with the increase of the Poisson process share. In the extreme situation when $\psi \leq 20\%$, the large $\hat{\mu}$ leads to an erroneous small expected number of SFP events in the observation time interval. Indeed, a mixture with Poisson process events only has $\mu = \infty$. Additionally to the bias problem, the estimator $\hat{\mu}$ also has a large variance when the SFP process has a small number of events.

We believe that the poor performance of the EM algorithm estimator $\hat{\mu}$ is related to the calculation of the expected value of the

likelihood function. This calculation was done using approximations to deal with the unknown events' labels, which directly influences the calculus of the SFP stochastic intensity function. This influence has less impact for the λ_p , since the Poisson process intensity is deterministic and fixed during the entire interval.




Figure 9: Boxplots of $\Delta(\hat{\mu}, \mu)$ for the improved estimator $\hat{\mu}$ according to $(100 - \psi)$ and n .

As μ is the median inter-event time in a pure SFP process, a simple and robust estimator in this pure SFP situation is the empirical median of the intervals $t_{i+1} - t_i$. Our alternative estimator for μ deletes some carefully selected events from the mixture, reducing the dataset to a pure SFP process and, then, taking the median of the inter-event times of the remaining events. More specifically, conditioned on the well behaved $\hat{\lambda}_p$ estimate, we generate pseudo events $u_1 < u_2 < \dots < u_m$ coming from a homogeneous Poisson process and within the time interval $(0, t_n)$. Sequentially define

$$t_k^* = \arg \min_{t_i \in S_k} |u_k - t_i|$$

with $S_k = T_k \cap R_k$ where

$$T_k = \{t_1, t_2, \dots, t_n\} - \{t_1^*, \dots, t_{k-1}^*\}$$

and $R_k = \{t_i : |u_k - t_i| < 2/\lambda_p\}$. This last constraint avoids the deletion to be entirely concentrated in bursty regions. We assume that the left over events in T_m constitute a realization of a pure SFP process and we use their median inter-event time as an estimator of μ . As this is clearly affected by the randomly deleted events t_k^* , we repeat this procedure many times and average the results to end up with a final estimate, which we will denote by $\hat{\mu}$.

The results obtained with the new estimator of μ can be visualized in Figure 9. Its estimation error is significantly smaller, with an underestimation of μ only when the PP component is dominant. This is expected because when the SFP component has a small percentage of events, its corresponding estimate is highly variable. In this case, there will be a large number of supposedly Poisson points deleted, remaining few SFP events to estimate the μ parameter, implying a high instability.

5. CLASSIFICATION TEST

When analysing a point process dataset, a preliminary analysis should test if a simpler point process, comprised either by a pure SFP or a pure Poisson process, fits the observed data as well as the more complex mixture model. Let $\theta = (\lambda_p, \mu)$ and the unconstrained parameter space be $\Theta = [0, \infty]^2$. We used the maximum likelihood ratio test statistic R of $H_1 : \theta \in \Theta$ against the null hypothesis $H_0 : \theta \in \Theta_0$ where, alternatively, we consider either $\theta \in \Theta_0 = (0, \infty) \times \{\infty\}$ or $\theta \in \Theta_0 = \{0\} \times (0, \infty)$ to represent

the pure Poisson and the pure SFP processes, respectively. Then

$$R = 2 \times \left(\max_{\theta \in \Theta} \ell(\theta) - \max_{\theta \in \Theta_0} \ell(\theta_0) \right)$$

where the log-likelihood $\ell(\theta)$ is given in (3). As a guide, we used a threshold $\alpha = 0.05$ to deem the test significant. As a practical issue, since taking the median inter-event time μ of the SFP process equal to ∞ is not numerically feasible, we set it equal to the length of the observed total time interval.

As there is one free parameter in each case, one could expect that the usual asymptotic distribution of $2 \log(R)$ should follow a chi-square distribution with one degree of freedom. However, this classic result requires several strict assumptions about the stochastic nature of the data, foremost the independence of the observations, which is not the situation in our model. Therefore, to check the accuracy of this asymptotic distribution to gauge the test-based decisions, we carried out 2000 additional Monte Carlo simulations, half of them following a pure Poisson process, the other half following a pure SFP. Adding these pure cases to those of the mixed cases at different percentage compositions described previously, we calculated the test p-values ϕ_p and ϕ_s based on the usual chi-square distribution with one degree of freedom. Namely, with \mathbb{F} being the cumulative distribution function of the chi-square distribution with one degree of freedom, we have

$$\phi_p = 1 - \mathbb{F}(R) = 1 - \mathbb{F} \left(2 \max_{\lambda_p, \mu} \ell(\lambda_p, \mu) - 2 \max_{\lambda_p} \ell(\lambda_p, \infty) \right) \quad (15)$$

and

$$\phi_s = 1 - \mathbb{F}(R) = 1 - \mathbb{F} \left(2 \max_{\lambda_p, \mu} \ell(\lambda_p, \mu) - 2 \max_{\mu} \ell(0, \mu) \right). \quad (16)$$

A figure showing the p-values ϕ_p and ϕ_s for all simulated point processes, pure or mixed, can be seen in the extended version of this paper. Summarizing the results, we have all ϕ_p practically collapsed to zero when ψ is large. The test will reject the null hypothesis that the process is a pure Poisson process, which is the correct decision. Indeed, this correct decision is taken in virtually all cases until $\psi \geq 40\%$. The test still correctly rejects the pure Poisson in all cases where $\psi > 20\%$ except when the number of events is very small. Only when the $\psi = 10\%$ or $\psi = 0\%$ (and, therefore, it is pure Poisson process) the p-value distribution clearly shifts upward and starts accepting the null hypothesis frequently. This is exactly the expected and desired behavior for our test statistic. The behavior of ϕ_s is analogous to that of ϕ_p . We also analysed the joint distribution of (ϕ_p, ϕ_s) . The two tests practically never accept both null hypothesis, the pure Poisson and pure SFP processes. Either one or other pure process is accepted or else both pure processes are rejected, indicating a mixed process.

6. FITTING AND CHARACTERIZATION

We used eight datasets split into three groups. The first one contains the comments on topics of several web services: the discussion forums *AskMe*, *MetaFilter*, and *MetaTalk* and the collaborative recommendation systems *Digg* and *Reddit*. The second group contains user communication events: e-mail exchange (*Enron*) and hashtag-based chat (*Twitter*). The fourth group is composed by user reviews and recommendations of restaurants in a collaborative platform (*Yelp*). In total, we analysed 18, 685, 678 events.

The *AskMe*, *MetaFilter*² and *MetaTalk* datasets were made avail-

²<http://stuff.metafilter.com/infodump/> - Accessed in September, 2013

able by the *Metafilter Infodump Project*². The *Digg*³ dataset was temporarily available on the web and was downloaded by the authors. The *Enron*⁴ data were obtained through the *CALO Project (A Cognitive Assistant that Learns and Organizes)* of Carnegie Mellon University. The *Yelp*⁵ data were available during the *Yelp Dataset Challenge*. The *Reddit* and *Twitter* datasets were collected using their respective APIs. All datasets have time scale where the unit is the second except for *Yelp*, which has the time scale measured in days, a more natural scale for this kind of evaluation review service.

For all databases, each RSE is a sequence of events timestamps and the event varies according to the dataset. For the *Enron* dataset, the RSE is associated with individual users and the events are the incoming and outgoing e-mail timestamps. For the *Twitter* dataset, each RSE is associated with a hashtag and the events are the tweet timestamps mentioning that hashtag. For the *Yelp* dataset, each RSE is associated with a venue and the events are the reviews timestamps. For all other datasets, the RSE is a discussion topic and the events are composed by comments timestamps. As verified by [22], the rate at which comments arrive has a drastic decay after the topic leaves the forum main page. The average percentage of comments made before this inflection point varies from 85% to 95% and these represents the bulk of the topic life. As a safe cutoff point, we considered the 75% of the initial flow of comments in each forum topic.

Table 1 shows the number of RSEs in each database. It also shows the average number of events by dataset, as well as the minimum and maximum number of events. We applied our classification test from Section 5 and the table shows the percentage categorized as pure Poisson process, pure SFP, or mixed process. For all datasets, the p-values ϕ_p and ϕ_s have the expected behavior, leading us to believe in the efficacy of our classification test to separate out the models in real databases in addition to their excellent performance in the synthetic databases. This is a result that can be verified in the extended version of this paper in the Arxiv website.

A more visual and complete way to look at the *burstiness scale* ψ is in the histograms of Figure 10. The horizontal axis shows the expected percentage of the Poisson process component in the RSE given by $\hat{\lambda}_p/n$. The two extreme bars at the horizontal axis, at $\psi = 100$ and $\psi = 0$, have areas equal to the percentage of series classified as pure SFP and as pure Poisson, respectively. The middle bars represent the RSE classified as mixed point processes. *AskMe*, *MetaFilter*, *MetaTalk*, *Reddit*, *Yelp* have the composition where the three models, the two pure and the mixed one, appear with substantial amount. The Poisson process share of the mixed processes distributed over a large range, from close to zero to large percentages, reflecting the wide variety of series behavior.

Figure 11 shows the estimated pairs $(\log \hat{\lambda}_p, \log \hat{\mu})$ of each events stream classified as a mixed process. The logarithmic scale provides the correct scale to fit the asymptotic bivariate Gaussian distribution of the maximum likelihood estimator. Each point represents a RSE and they are colored according to the database name. Except by the *Twitter* dataset, all others have their estimator distribution approximately fitted by a bivariate Gaussian distribution with marginal mean, variance and correlation given in Table 1.

The correlation is negative in all databases, meaning that a large value of the Poisson process parameter (that is, a large $\hat{\lambda}_p$) tends to be followed by small values of the SFP component (that is, a

³<http://www.infochimps.com/datasets/diggcom-data-set> - Accessed in September, 2013

⁴<http://www.cs.cmu.edu/~enron/> - Accessed in September, 2013

⁵http://www.yelp.com/dataset_challenge - Accessed in August, 2014

Table 1: Description of the databases: number of series of events; minimum, average, and maximum number of events; classification test results; gaussian fit parameters.

Base	# of series	# of events per series			Hypothesis Test			Bivariate Gaussian		
		Min	Avg	Max	Mix	PP	SFP	$\log \lambda_P (\sigma^2_{\log \lambda_P})$	$\log \hat{\mu} (\sigma^2_{\log \mu})$	$\rho(\log \lambda_P, \log \mu)$
AskMe	490	74	99.30	699	333 (67.96%)	43 (8.78%)	114 (23.26%)	-6.91 (0.98)	4.77 (0.39)	-0.40
Digg	974	39	90.41	296	353 (36.24%)	2 (0.21%)	619 (63.55%)	-8.08 (0.83)	4.4 (0.38)	-0.11
Enron	145	55	1,541.35	14258	106 (73.1%)	0 (0%)	39 (26.9%)	-11.56 (0.62)	8.18 (0.57)	-0.28
MetaFilter	8243	72	131.10	4148	5625 (68.24%)	1279 (15.52%)	1339 (16.24%)	-6.76 (0.94)	4.78 (0.37)	-0.39
MetaTalk	2460	73	151.92	2714	1691 (68.74%)	271 (11.02%)	498 (20.24%)	-7.31 (1.08)	5.23 (0.55)	-0.57
Reddit	102	37	535.43	4706	58 (56.86%)	21 (20.59%)	23 (22.55%)	-6.12 (1.1)	3.26 (3.48)	-0.85
Twitter	17088	50	969.68	8564	15913 (93.12%)	72 (0.42%)	1103 (6.46%)	-10.01 (0.31)	6.82 (5.26)	-0.66
Yelp	1929	50	127.84	1646	774 (40.12%)	927 (48.06%)	228 (11.82%)	-3.79 (0.38)	2.28 (0.38)	-0.22

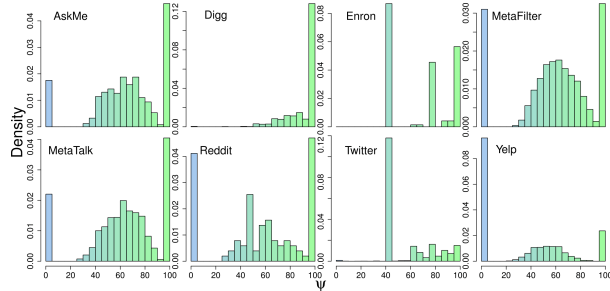


Figure 10: The burstiness scale ψ in each dataset. RSEs range from being fully Poissonian ($\psi = 0$) to completely bursty ($\psi = 100$).

small $\hat{\mu}$, implying a short median inter-event time between the SFP events). Not only each database has a negative correlation between the mixture parameters, they also occupy a distinct region along a NorthWest-SouthEast gradient. Starting from the upper left corner, we have the Enron email cloud, exhibiting a low average $\hat{\lambda}_P$ and a jointly high $\hat{\mu}$. Descending the gradient, we find the less compactly shaped Twitter point cloud. In the $(-8, -6) \times (4, 6)$ region we find the forums (AskMe, MetaFilter, MetaTalk). Slightly shifted to the left and further below (within the $(-10, -7) \times (3, 5)$ region), we find the two collaborative recommendation systems (Reddit and Digg). Finally, in the lower right corner, we have the Yelp random series estimates.

In this way, our model has been able to spread out the different databases in the space composed by the two component processes parameters. Different communication services lives in a distinctive location in this mathematical geography.

7. GOODNESS OF FIT

Figure 12 shows a goodness of fit statistic for the RSE classified as a mixed process. After obtaining the $\hat{\lambda}_P$ and $\hat{\mu}$ estimates, we disentangled the two processes using the Monte Carlo simulation procedure described in Section 4.2. The separated out events were then used to calculate the statistics shown in the two histograms. The plot in Figure 12a is the determination coefficient R^2 from the linear regression of the events cumulative number $N(t)$ versus t , which should be approximately a straight line under the Poisson process.

In Figure 12b we show the R^2 from a linear regression with the SFP-labelled events. We take the inter-event times sample and build the empirical cumulative distribution function $\mathbb{F}(t)$ leading to the odds-ratio function $OR(t) = \mathbb{F}(t)/(1 - \mathbb{F}(t))$. This function should be approximately a straight line if the SFP process hypothesis is valid (more details in [19]).

Indeed, the two histograms of Figure 12 show very high concen-

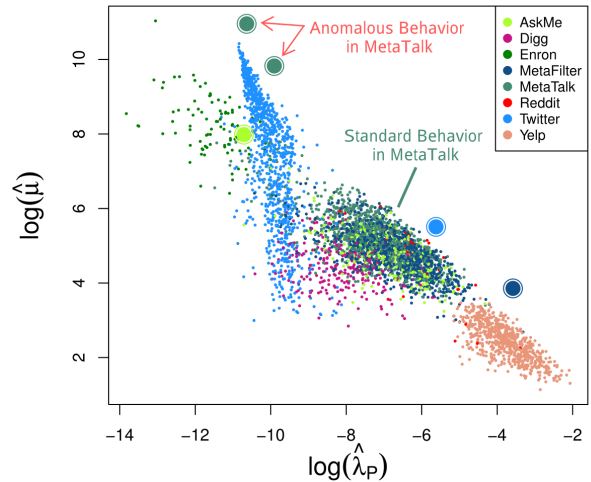


Figure 11: Estimates $(\hat{\lambda}_P, \hat{\mu})$ for all events streams from the eight databases (logarithmic scale). A few anomalous time series are highlighted as large dots.

tration of the R^2 statistics close to the maximum value of 1 for the collection of RSEs. This provides evidence that our disentangling procedure of the mixed process into two components is able to create two processes that fit the characteristics of a Poisson process and a SFP process.




Figure 12: Goodness of fit of mixed series.

8. COMPARISON WITH HAWKES PROCESS

An alternative process to our model, is the Hawkes point process [3], which has conditional intensity defined by

$$\lambda(t|\mathcal{H}_t) = \lambda_p + \sum_{t_i < t} K(t - t_i) \quad (17)$$

where $K(x) > 0$ is called the kernel function. As our *BuSca*

model, the Hawkes process allows the successive events to interact with each other. However, there are two important differences between them. In Hawkes, every single event excites the process increasing the chance of additional events immediately after, while only some of these incoming events induces process excitement in BuSca. Depending on the value of ψ , only a fraction of the events lead to an increase on $\lambda(t|\mathcal{H}_t)$. The second difference is the need to specify a functional form for the kernel $K(x)$, common choices being $K(x)$ with an exponential decay or a power law decay.

We compared our model with the alternative Hawkes process using the 31431 events time series of all databases we analysed. We fitted both, *BuSca* and the Hawkes process, at each time series separately by maximum likelihood, and evaluated the resulting Akaike information criterion. The Hawkes process was fitted with the exponential kernel implying on a three parameter model while *BuSca* requires only two. The result was: only four out of 31431 RSEs had their Hawkes AIC smaller. Therefore, for practically all time series we considered, our model fits better the data although requiring less parameters.

To understand better its relative failure, we looked at the R^2 of the fitted Hawkes model for each time series (see [15] for the R^2 calculation in the Hawkes model). We studied the R^2 distribution conditioned on the value of ψ . Hawkes is able to fit reasonably well only when $\psi < 30\%$, that is, when the series of events is Poisson dominated. It has mixed results for $0.3 \leq \psi \leq 0.5$, and a poor fit when $\psi > 0.5$, exactly when the bursty periods are more prevalent.

9. APPLICATIONS

9.1 Anomaly detection

Within a given database, we saw empirically that the maximum likelihood estimator ($\log \hat{\lambda}_p$, $\log \hat{\mu}$) of the series of events follows approximately a bivariate Gaussian distribution, as indeed we see in Figure 11. This provides a direct score for an anomalous point time series based on the Mahalanobis distance distribution with threshold $\alpha = 0.01$ (see details in the extended version).

In Figure 11 we highlighted 5 anomalous points found by our procedure to illustrate its usefulness. The first one correspond to the topic #219940 of the *AskMe* dataset, which has a very low value for $\hat{\lambda}_s$, compared to the other topics in this forum. This topic was initiated by a post about a lost dog and his owner asking for help. Figure 13a shows the cumulative number $N(t)$ of events up to time t , measured in days. Consistent with the standard behavior in this forum, there is an initial burst of events with users suggesting ways to locate the pet or sympathizing with the pet owner. This is followed by a Poissonian period of events arising at a constant rate. Occasional bursts of lower intensity are still present but eventually the topic reaches a very low rate. Typically, about $t = 12$ days, the topic would be considered dead and we would not see any additional activity. Further discussion from this point on would likely start a new topic. However, this was not what happened here. At $t = 15.9$, the time marked by the vertical red line, the long inactivity period is broken by a post from the pet owner mentioning that he received new and promising information about the dog whereabouts. Once again, he receives a cascade of suggestions and supporting messages. Before this flow of events decreases substantially, he posts at time $t = 17.9$ that the dog has been finally found. This is marked by the blue vertical line and it caused a new cascade of events congratulating the owners by the good news. This topic is anomalous with respect to the others in the *AskMe* database because, in fact, it contains three successive typical topics considered as a single one. The long inactivity period, in which the topic

was practically dead, led to a $\hat{\lambda}_p$ with a very low value, reflecting mathematically the anomaly in the content we just described.

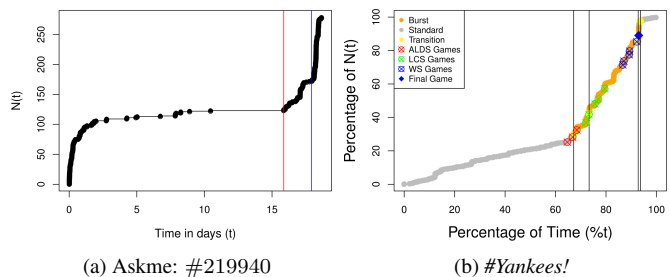


Figure 13: Applications: detection of anomaly and bursts.

In the *MetaTalk* database, the time series # 18067 and # 21900 deal with a unusual topic in this platform. They are reminders of the deadline for posting in an semestral event among the users and this prompted them to justify their lateness or make a comment about the event. This event is called *MeFiSwap* and it a way found by the users to share their favorite playlists. The first one occurred in the summer of 2009 and the second one in the winter of 2012. Being reminders, they do not add content, but refer to and promote other forum posts. What is anomalous in these two time series is the time they took to develop: 22.6 days (# 18067) and 10.9 (# 21900), while the average topic takes about 1.9 days. The pattern within their enlarged time scale is the same as the rest of the database. The behavior of these two cases is closer to the Twitter population, as can be seen in Figure 11.

The Twitter time series # 1088 was pinpointed as an anomaly due to his relatively large value of the Poisson component λ_p . It offered free tickets for certain cultural event. To qualify for the tickets, users should post something using the hashtag *iwantisatickets*. This triggered a cascade of associated posts that kept an approximately constant rate while it lasted.

Our final anomaly example is the time series # 65232 from *MetaFilter*. It was considered inappropriate and deleted from the forum by a moderator. The topic author suggested that grocery shopping should be exclusively a women’s chore because his wife had discovered many deals he was unable to find out. The subject was considered irrelevant and quickly prompted many criticisms.

9.2 Burst detection and identification

Another application developed is the detection and identification of burst periods in each individual time series. The idea is that a period with essentially no SFP activity should have the cumulative number of events $N(t)$ increasing at a constant and minimum slope approximately equal to λ_p . Periods with SFP activity would quickly increase this slope to some value $\lambda_p + c$. We explore this intuitive idea by segmenting optimally the $N(t)$ series.

We explain our method using Figure 13b, showing the history of the Twitter hashtag *Yankees!* spanning the regular and postseason periods in 2009. The detailed explanation is in the extended version of this paper. Briefly, we use the *Segmented Least Squares* algorithm from [10] to find an optimal segmentation, as shown in Figure 13b, together with the main New York Yankees games during playoffs (ALDS, LCS, and WS). In each segment, we calculate $\tau(s)$ which measures how much more bursty events we have in comparison with the underlying Poisson. The large $\tau(s)$, the more intense the burst in that s segment.

During the regular season, with posts coming essentially from the more enthusiastic fans, the behavior is completely dominated by a homogeneous Poisson process. The first segment found by the

algorithm starts during the October 7 week, at the first postseason games against Minnesota Twins (red crosses in Figure 13b). In this first playoff segment, we have $\tau(s) \approx 3$, or four times the standard regular behavior. The LCS games start an augmented burst until November 4, when the New York Yankees defeated the Philadelphia Phillies in a final game. This last game generated a very short burst marked by the blue diamond, with $\tau(s) \approx 64$. After this explosive period, the series resume to the usual standard behavior.

Analysing the $\tau(s)$ statistical distribution for each database separately, we found that they are well fitted by a heavy tailed probability distribution, a finding that is consistent with previous studies of cascade events [1].

10. CONCLUSIONS

In this paper, we proposed the *Burstiness Scale (BuSca)* model, which views each random series of events (RSEs) as a mix of two independent process: a *Poissonian* and a self-exciting one. We presented and validated a particular and parsimonious case of *BuSca*, where the Poissonian process is given by a homogeneous Poisson process (PP) and the self-exciting process is given by a Self-Feeding Process (SFP) [20]. When constructed in this way, *BuSca* requires two parameters to characterize RSEs, one for the PP and another for the SFP. We validated our approach by analyzing eight diverse and large datasets containing real RSEs seen in Twitter, Yelp, e-mail conversations, Digg, and online forums. We also proposed a method that uses the *BuSca* model to disentangle events related to routine and constant behavior (Poissonian) from bursty and trendy ones (self-exciting). Moreover, from the two parameters of *BuSca*, we can calculate the burstiness scale parameter ψ , which represents how much of the RSE is due to bursty and viral effects. We showed that these two parameters, together with our proposed burstiness scale, is a sparing way to accurately characterize random series of events, which and, consequently, may leverage several applications, such as monitoring systems, anomaly detection methods, flow predictors, among others.

Acknowledgments

This work was funded by author's individual grants from CAPES, CNPq, and Fapemig.

11. REFERENCES

- [1] A.-L. Barabási. The origin of bursts and heavy tails in human dynamics. *Nature*, 435(7039):207–211, may 2005.
- [2] J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun. Internet Traffic Tends Toward Poisson and Independent as the Load Increases. pages 83–109. 2003.
- [3] R. Crane and D. Sornette. Robust dynamic classes revealed by measuring the response function of a social system. *Proceedings of the National Academy of Sciences*, 105(41):15649–15653, oct 2008.
- [4] A. Ferraz Costa, Y. Yamaguchi, A. Juci Machado Traina, C. Traina, and C. Faloutsos. Rsc: Mining and modeling temporal activity in social media. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '15*, pages 269–278, New York, New York, USA, 2015. ACM Press.
- [5] F. Figueiredo, J. M. Almeida, Y. Matsubara, B. Ribeiro, and C. Faloutsos. Revisit Behavior in Social Media: The Phoenix-R Model and Discoveries. pages 386–401. 2014.
- [6] F. A. Haight. *Handbook of the Poisson distribution [by] Frank A. Haight*. Wiley New York., 1967.
- [7] H. Jiang and C. Dovrolis. Why is the Internet Traffic Bursty in Short Time Scales? In *Proceedings of the 2005 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '05)*, pages 241–252, 2005.
- [8] T. Karagiannis, M. Molle, M. Faloutsos, and A. Broido. A nonstationary poisson view of internet traffic. In *IEEE INFOCOM 2004*, volume 3, pages 1558–1569. IEEE.
- [9] J. Kleinberg. Bursty and hierarchical structure in streams. In *Proceedings of the eighth ACM SIGKDD*, KDD '02, pages 91–101, New York, NY, USA, 2002. ACM.
- [10] J. Kleinberg and E. Tardos. *Algorithm design*. Pearson Education, 2006.
- [11] J. Lehmann, B. Gonçalves, J. J. Ramasco, and C. Cattuto. Dynamical classes of collective attention in twitter. In *Proceedings of the 21st international conference on World Wide Web - WWW '12*, page 251, New York, New York, USA, 2012. ACM Press.
- [12] R. D. Malmgren, J. M. Hofman, L. A. Amaral, and D. J. Watts. Characterizing individual communication patterns. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '09*, page 607, New York, New York, USA, 2009. ACM Press.
- [13] Y. Matsubara, Y. Sakurai, B. A. Prakash, L. Li, and C. Faloutsos. Rise and fall patterns of information diffusion. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '12*, page 6, New York, New York, USA, 2012. ACM Press.
- [14] J. G. Oliveira and A.-L. Barabasi. Human dynamics: Darwin and Einstein correspondence patterns. *Nature*, 437(7063):1251, 2005.
- [15] R. Peng. Multi-dimensional point process models in r. *Journal of Statistical Software*, 8(1):1–27, 2003.
- [16] J. C. L. Pinto, T. Chahed, and E. Altman. Trend detection in social networks using Hawkes processes. In *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015 - ASONAM '15*, pages 1441–1448, New York, New York, USA, 2015. ACM Press.
- [17] D. L. Snyder and M. I. Miller. *Random Point Processes in Time and Space*. Springer Texts in Electrical Engineering. Springer New York, New York, NY, 1991.
- [18] D. Vallet, S. Berkovsky, S. Ardon, A. Mahanti, and M. A. Kafaar. Characterizing and Predicting Viral-and-Popular Video Content. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*. CIKM '15, pages 1591–1600, New York, NY, USA, 2015. ACM.
- [19] P. O. S. Vaz de Melo, C. Faloutsos, R. Assunção, R. Alves, and A. A. Loureiro. Universal and Distinct Properties of Communication Dynamics: How to Generate Realistic Inter-event Times. *ACM Transactions on Knowledge Discovery in Data*, 2015.
- [20] P. O. S. Vaz de Melo, C. Faloutsos, R. Assuncao, and A. A. F. Loureiro. The Self-Feeding Process: A Unifying Model for Communication Dynamics in the Web. In *WWW '13: 22nd International World Wide Web Conference*, 2013.
- [21] A. Vazquez, J. G. Oliveira, Z. Dezso, K.-I. Goh, I. Kondor, and A.-L. Barabasi. Modeling bursts and heavy tails in human dynamics. *Phys Rev E Stat Nonlin Soft Matter Phys*, 73:36127, 2006.
- [22] C. Wang, M. Ye, and B. A. Huberman. From user comments to on-line conversations. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '12*, page 244, New York, New York, USA, 2012. ACM Press.
- [23] S. Wang, Z. Yan, X. Hu, P. S. Yu, and Z. Li. Burst Time Prediction in Cascades, 2015.
- [24] J. Yang and J. Leskovec. Patterns of temporal variation in online media. In *Proceedings of the fourth ACM international conference on Web search and data mining - WSDM '11*, page 177, New York, New York, USA, 2011. ACM Press.
- [25] S.-h. Yang and H. Zha. Mixture of Mutually Exciting Processes for Viral Diffusion. In S. Dasgupta and D. Mcallester, editors, *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, volume 28, pages 1–9. JMLR Workshop and Conference Proceedings, 2013.
- [26] H. Yu, L. Xie, and S. Sanner. The Lifecycle of a Youtube Video: Phases, Content and Popularity, 2015.
- [27] Q. Zhao, M. A. Erdogdu, H. Y. He, A. Rajaraman, and J. Leskovec. SEISMIC: A Self-Exciting Point Process Model for Predicting Tweet Popularity. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '15*, pages 1513–1522, New York, New York, USA, 2015. ACM Press.