

RESEARCH ARTICLE

Predicting Relationship Labels and Individual Personality Traits From Telecommunication History in Social Networks Using Hawkes Processes

MATEUSZ NUREK¹, RADOSŁAW MICHALSKI¹, (Member, IEEE), OMAR LIZARDO², AND MARIAN-ANDREI RIZOIU³

¹Department of Artificial Intelligence, Wrocław University of Science and Technology, 50-370 Wrocław, Poland

²Department of Sociology, University of California at Los Angeles, Los Angeles, CA 90095, USA

³Data Science Institute, University of Technology Sydney, Sydney, NSW 2007, Australia

Corresponding author: Mateusz Nurek (mateusz.nurek@pwr.edu.pl)

This work was supported in part by the National Science Centre, Poland, under Grant 2021/41/B/HS6/02798; in part by the Commonwealth of Australia (represented by the Defence Science and Technology Group) through a Defence Science Partnerships Agreement; in part by the NetSense Study through the U.S. National Science Foundation [Computer and Information Science and Engineering (CISE)] under Grant 0968529; and in part by the NetHealth Study through the U.S. National Institutes of Health (NIH) under Grant R01HL117757.

ABSTRACT Mobile phones contain a wealth of private information, so we try to keep them secure. We provide large-scale evidence that the psychological profiles of individuals and their relations with their peers can be predicted from seemingly anonymous communication traces—calling and texting logs that service providers routinely collect. Based on two extensive longitudinal studies containing more than 900 college students, we use point process modeling to describe communication patterns. We automatically predict the peer relationship type and temporal dynamics, and assess user personality based on the modeling. For some personality traits, the results are comparable to the gold-standard performances obtained from survey self-report data. Findings illustrate how information usually residing outside the control of individuals can be used to reconstruct sensitive information.

INDEX TERMS Big5, call series modeling, Hawkes process, inferring personality traits, inferring relation status, relationship dynamic.

I. INTRODUCTION

It is a well-known fact that we give away personal information whenever we act and interact online [49]. This has been repeatedly shown for online social media platforms such as Facebook [27], [62], Twitter [33], and even knowledge creation sites like Wikipedia [47]. More specifically, people's personality profiles can be predicted from the activity

The associate editor coordinating the review of this manuscript and approving it for publication was Giacomo Fiumara¹.

traces they leave behind in personal websites [34], blogs [61], Twitter messages [22] or Facebook profiles [27].

Recently, Stachl et al. [52] showed that personality profiles could be estimated using information collected from users' smartphones. This includes digital traces of communication and social behavior, music consumption, app usage, mobility, overall phone activity, and day- and nighttime activity. However, obtaining such data requires access to the user's social media activity or phones, which might provide specific degrees of safety to the user who controls the privacy settings. In this work, we investigate what can be learned about

users using data normally outside their control, e.g., from communication patterns with their peers. We find that even seemingly anonymous information can be used to predict, with a high level of accuracy both individual characteristics and properties of their social relationships.

This paper addresses three open questions concerning modeling and learning from call and text data traces. The first question relates to modeling the communication patterns between individuals. It is known that human communication is bursty and that it exhibits a long-tail distribution of inter-event times [19], [32], [58]. Hawkes point processes have been successfully applied to model other bursty phenomena, such as information diffusion [37] or neuronal firing patterns in the human brain [21]. The question is **can we model the call patterns between individuals using Hawkes point processes?** The second question relates to learning about relationships between individuals using call patterns. Hawkes processes fitted on event series have been shown helpful in predicting the final popularity of online items [37], and even differentiating between controversial and authoritative news sources [25]. The question is, therefore, **can we differentiate the relationship type between two individuals, given their call and texting timing series fitted using a Hawkes model?** The third question relates to inferring users' personality traits from their call series. While there exist several prior work learning user personality traits using mobile phone data [16], [38], [52], these usually rely on behavioral information collected via sensor and log data from smartphones, which requires access to the users' phones. The question is **can we use the Hawkes model outputs, fitted on the call series involving a given user, to infer their personality traits?**

A. THE $\text{NET}_{\text{SENSE}}$ AND $\text{NET}_{\text{HEALTH}}$ STUDIES

We address the above open questions using two large datasets – $\text{NET}_{\text{SENSE}}$ and $\text{NET}_{\text{HEALTH}}$ – obtained from the NetSense and NetHealth longitudinal social network studies conducted on a student population at the University of Notre Dame [43], [53]. The datasets contain digital trace timestamped records of the calling and texting information for about 200 and 700 students, respectively, who also filled in periodic surveys administered at the start and end of each semester for two and three years in each respective study. We answer the first question by fitting the parameters of a Hawkes point process to the series of communication events – phone calls or texts – occurring between each pair of students. Using goodness of fit tests, we show that Hawkes with a power-law decaying kernel function generalizes better to unseen data than the exponential kernel.

To answer the second question, we use the student surveys in which they provided labels for their most salient (top-20) contacts in periodic ego-network surveys (e.g., *family*, *friend*, *significant other* and the like). Based on the labels given to relationships by students in subsequent surveys, we first split the pairwise relations in both datasets into six categories: family-relaxing, family-stable, friendship-

relaxing, friendship-stable, friendship-strengthening and romantic-relaxing. Next, we characterize each pairwise relationship in our dataset using the Hawkes model's fitted parameters and secondary quantities. We show that using the Hawkes descriptors, off-the-shelf classifiers significantly outperform autoregressive (ARIMA [31]) baselines, achieving a macro F1-score of 0.24 for $\text{NET}_{\text{SENSE}}$ and 0.22 for $\text{NET}_{\text{HEALTH}}$, with the best-identified categories being family-stable, friendship-stable, and friendship-relaxing. We obtain the best prediction results when concatenating the autoregressive and Hawkes features (F1-score of 0.27 and 0.23 on $\text{NET}_{\text{SENSE}}$ and $\text{NET}_{\text{HEALTH}}$, respectively).

We answer the third question by applying a novel joint modeling of Hawkes processes [25] to model jointly the call series of each student participating in the $\text{NET}_{\text{SENSE}}$ and $\text{NET}_{\text{HEALTH}}$ studies. We describe users using features built on their activity patterns and we use off-the-shelf regressors to predict individual Big5 personality traits [15]. The ground truth is the Big Five Inventory [20], a personality assessment survey filled in by each student. We compare our method with two baselines. The first baseline [38] uses a wide range of complex features (including autoregressive) based on call history. The second baseline uses 145 features ($\text{NET}_{\text{SENSE}}$) and 2212 features ($\text{NET}_{\text{HEALTH}}$) self-reported in the surveys: grades, health, happiness, activity, book reading, and club membership. These are highly personal and sensitive features, and have been shown to be the upper bound of automatic personality traits prediction [40]. We find our method to outperform the first baseline and, surprisingly, very competitive with respect to the prediction upper-bound. This indicates that the Hawkes-modelled call activities embed a surprising amount of personal information.

The main contributions of this paper are as follows:

- We show that a Hawkes point process with a power-law decaying kernel can model the phone call contact series between individuals.
- We use the fitted parameters of a Hawkes model to distinguish between types of relations (such as family, friendship, or romantic).
- We show that the call activity (modeled using Hawkes point processes) might be as predictive of the user's Big5 psychometric traits as the user filled-in questionnaire.

B. THE ETHICS OF PERSONALITY PROFILING

Personality profiling—particularly social media-inferred personality traits—is sometimes seen as Pandora's box. On the one hand, personality dispositions are associated with happiness, physical and psychological health, the quality of relationships with peers, family, and romantic others, as well as and community involvement [42], criminal activity, and political ideology at a social, institutional level [42]. Personality traits are also predictive of three critical life outcomes: mortality, divorce, and occupational attainment [48]. Such research shows the positive aspects of personality profiling research: one could build systems to prevent and improve individuals' mental health issues or their

relationship with the community. On the other hand, it was also shown that persuasive messaging is more effective when tailored for individuals’ psychological characteristics [36] and that the same processes used to infer personality traits from social data can leak sensitive information such as ethnicity, political and sexual orientation [27]. While social media privacy settings could (at least theoretically) bring some of the data back under the user’s control, our research shows that personality traits can be inferred from data sources entirely outside the user’s control. This work adds to the understanding of what can be achieved with call logs data and advocates creating policy regulating its usage. The latter is increasingly important, as more and more calls are being made using outside the traditional communication networks and onto online messaging platforms such as WhatsApp, which are currently not nationally regulated.

II. PREREQUISITES: HAWKES PROCESSES

In this section, we briefly review the theoretical prerequisites concerning modeling event series using Hawkes point processes.

A. EVENT SERIES AND POINT PROCESSES

An event is a tuple (timestamp, event features), where the *timestamp* sits along the non-negative time axis, and the *event features* are any descriptors related to the event. For example, an event can be the reception of a phone call [44], [56], and the features could be the caller id, length of the call, or whether it was answered or not. An event series is a sequence of events t_1, t_2, \dots , where t_i are the event timestamps of the i^{th} event, relative to the first event ($t_0 = 0$). For ease of notation, in this paper, we use t_i to denote both the event timestamp and the event itself. A point process is a random process whose realizations consist of event series. We denote an event series observed up to time T as $\mathcal{H}(T) = \{t_0, t_1, \dots\}$.

B. THE HAWKES PROCESSES

Hawkes processes [17] are a type of point process with the self-exciting property, i.e., the occurrence of past events increases the likelihood of future events. This results in the cluster property of the Hawkes property [18], which states that events modeled by Hawkes appear to be grouped in time. This latter property makes Hawkes processes desirable to model human interaction activity, which is known to follow a bursty pattern [63]. An alternative approach for modeling interactions are Wold point processes defined through a Markovian transition probability distribution on the inter-event times [11]. The occurrence of events in a Hawkes process is controlled by the event intensity function:

$$\lambda(t | \mathcal{H}(T)) = \mu(t) + \sum_{t_i < t} \phi(t - t_i) \tag{1}$$

where $\mu(t)$ is the background intensity function and $\phi : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is a kernel function capturing the decaying influence of a past event. $\mu(t)$ is the immigrant intensity (i.e. the volumes of events that come from outside the system).

Here, it is constant for each call series, i.e., $\mu(t) = \mu$. Two widely adopted parametric forms for the kernel function ϕ include the exponential function $\phi_{EXP}(t) = \kappa\theta e^{-\theta t}$ and the power-law function $\phi_{PL}(t) = \kappa(t + c)^{-(1+\theta)}$.

C. THE BRANCHING FACTOR

n^* is defined as the expected number of events directly spawned by a single event, i.e., $n^* = \int_0^\infty \phi(\tau) d\tau$. n^* defines the *regime* of the Hawkes process: when $n^* > 1$, each event generates more than one event; the process is expected to generate an infinite number of events. When $n^* < 1$, the process generates a finite number of events, and it is expected to die out.

D. PARAMETER ESTIMATION

We estimate the parameters of the Hawkes process by maximizing the log-likelihood function for point processes [9]. In a nutshell, the function quantifies the probability that the model generated the observed data. Higher probabilities indicate better fit models. Formally,

$$\mathcal{L}(\Theta | \mathcal{H}(T)) = \sum_{t_j \in \mathcal{H}(T)} \log \lambda(t_j | \mathcal{H}(T)) - \int_0^T \lambda(\tau | \mathcal{H}(T)) d\tau \tag{2}$$

E. JOINT MODELING OF EVENT SERIES

When analyzing the event series relating to a single entity (say, all the phone call series generated by the same individual), it is desirable to account for the multiple event series simultaneously. Kong et al. [25] proposed to jointly model a group of retweet cascades with a shared Hawkes process model by summing the log-likelihood functions of individual series. In *Section Inferring Psychometric Traits*, we jointly model the call sequences initiated by the same user, and we link the learned models to users’ psychometric traits.

III. DATA AND METHODS

In this section, we first introduce $\text{NET}_{\text{SENSE}}$ and $\text{NET}_{\text{HEALTH}}$, the two datasets used in this work (*Section Datasets*). Next, we build the relationship labels (*Section Relationships*), we profile the datasets (*Section Datasets profiling*), and we fit Hawkes processes to telecommunication data (*Section Analyzing call patterns using Hawkes*).

A. DATASETS

This research uses two datasets: $\text{NET}_{\text{SENSE}}$ [53] and $\text{NET}_{\text{HEALTH}}$ [29]. The $\text{NET}_{\text{SENSE}}$ project lasted for two and half years and gathered metadata regarding phone calls, demographic, and networking information about college students enrolled in the study during 2011–2013. The $\text{NET}_{\text{SENSE}}$ study was followed by $\text{NET}_{\text{HEALTH}}$, which also recorded metadata on mobile phone activities of Notre Dame University students but over a longer time – four and a half years from 2015 to 2019. Both datasets consist of two parts: the calling and texting activity and the student surveys.

1) ETHICS

All methods were carried out following relevant guidelines and regulations. The Institutional Review Board (IRB) of the University of Notre Dame has reviewed the $\text{NET}_{\text{SENSE}}$ study and has approved it - the IRB Number is FWA 00002462. The observational study $\text{NET}_{\text{HEALTH}}$ was also approved by the University of Notre Dame's IRB after a full board review under protocol 17-05-3912. All participants provided written informed consent prior to taking part in both studies.

2) CALLING AND TEXTING ACTIVITY

The $\text{NET}_{\text{SENSE}}$ and $\text{NET}_{\text{HEALTH}}$ datasets record all calls and texts from the students' phones enrolled in the study, both outgoing and incoming. The dataset records the caller id, the receiver id, the timestamp, and the call duration for each phone call. For text messages, it records the sender, the receiver, and the timestamp. $\text{NET}_{\text{HEALTH}}$ also records metadata about WhatsApp and iMessage communication. Where both the sender (or caller) and the receiver are students enrolled in the program, the call (or text) will be recorded twice – once for the caller and once for the receiver. $\text{NET}_{\text{SENSE}}$ and $\text{NET}_{\text{HEALTH}}$ also record calls and texts from people outside the study (such as family and friends).

3) THE SURVEYS

The datasets include detailed data on the students since each participating student was surveyed every term about their interests, opinions, or relationships with others. The students have been surveyed six times for $\text{NET}_{\text{SENSE}}$ and eight times for $\text{NET}_{\text{HEALTH}}$. This paper leverages two aspects of these surveys: how students describe their relationships with others and how they represent themselves. For the former, participants labeled the relationship with their peers using descriptors such as a friend, family, significant other, co-worker, other, or not labeled at all. For the latter, the students provided information about their hobbies, activities, well-being, grades, weight and height, health condition, and the number of books they read. We use this information in [Section *Inferring Psychometric Traits*](#) to build a baseline for predicting personality traits.

4) THE BIG5 PSYCHOMETRIC QUESTIONNAIRE

For both datasets and with each survey, students provided answers to forty-four questions from the Big Five Inventory [20]. Students answered the questions on whether they would describe themselves as *talkative*, *curious about many different things*, *tending to be quiet*, among others. Their answers can map each student as a point in the big five personality traits space [15], where each traits (openness, conscientiousness, extraversion, agreeableness, neuroticism) is represented as a numeric value between one and five.

B. RELATIONSHIPS

Research in Sociology, Psychology, Communication, and related fields show that social network relations are inherently

TABLE 1. The rules for categorizing relationships used in this study and the number of obtained instances for each rule.

Previous relation	New relation	Label	$\text{NET}_{\text{SENSE}}$	$\text{NET}_{\text{HEALTH}}$
friend friend friend	sibling parent other family	family- relaxing	115	313
parent sibling	parent sibling	family- stable	363	422
friend friend friend	coworker other acquaintance	friendship- relaxing	569	1709
friend	friend	friendship- stable	675	988
other coworker acquaintance	friend friend friend	friendship- strengthening	41	74
significant other significant other	other friend	romantic- relaxing	16	18

dynamic – they change in strength, quality, and character over time [12], [14], [45], [54], [55]. Strong ties can become weak [4], and weak ties can become strong [13]. Contacts may transition from initial acquaintances or workmates to being close friends [12], [60]. In the same way, geographical separation can change the character of ties, with close friends devolving into casual ties. The fundamental message of this previous work is that people's relations with their peers evolve.

To capture these dynamics, the $\text{NET}_{\text{SENSE}}$ and $\text{NET}_{\text{HEALTH}}$ projects queried students about their personal relationships with their top-twenty contacts six times during the respective study periods. In particular, each participant was asked to provide a relationship *label*, characterizing each tie as a friend, family (e.g., parent, sibling, other kin), co-worker, acquaintance, or significant other. We discriminate between three types of temporal dynamics for relationships: stable, strengthening, and relaxing. We denote a relation as stable when the student labels it identically across all surveys. When the relationship transitions between categories, we annotate it with the type and direction of the development, following the rules listed in Table 1. We consider a *friend* as a stronger type of relation than a *acquaintance*; therefore we label a transition *acquaintance* \rightarrow *friend* as “friendship-strengthening”. Similarly, we denote the transition *significant other* \rightarrow *friend* as “romantic-relaxing”. One type of transition that might appear peculiar is the “family-relaxing”, where the relation is initially labeled as *friend*, and later on, it becomes *sibling*, *parent* or *other family*. The reader should remember that the subjects of the studies were college students, for most of whom college is the first time they leave home and live independently. Therefore, it is not inconceivable that the relationship with the family members can go colder over time as they discover friends and college life. Finally, we discard the 13 relationships which have more than one type transition;

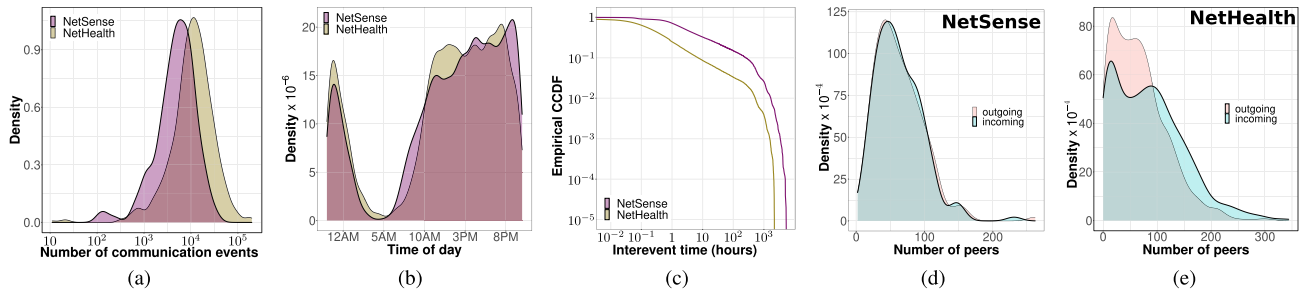


FIGURE 1. Exploratory analysis of the communication data for students participating in the NETSENSE and NETHEALTH studies at the University of Notre Dame. (a) density plot for the number of communication events for the students, (b) the distribution of phone activity through the course of the day, (c) the cumulative distribution function of the inter-event times in hours, (d) and (e) the number of incoming and outgoing communication peers for the students in NETSENSE and NETHEALTH studies, respectively, where there are at least twenty communication events with the peer during the analyzed period.

we consider these as likely input errors by the students (e.g., *friend* → *significant other* → *family* → *other*). We also discard the relationship classes with less than 15 instances to have enough instances in each class for building classifiers in [Section Classify relationship types](#). Table 1 shows each relationship type’s final volumes in each dataset.

C. DATASETS PROFILING

The NETSENSE dataset covers 178 students who filled out the surveys and whose mobile phone communication is available. NETSENSE records 7,575,864 phone interactions using the application installed on the students’ mobile phones. The NETHEALTH study covered 712 students, and it collected information about 60,486,565 phone interactions. In Fig. 1, we present the exploratory analysis of communication datasets. Fig. 1a shows the density of the number of communication events per student. Visibly, the distributions look very similar for both datasets, with most students recording between 5,000 and 20,000 communication events. The distribution is shifted towards higher values for NETHEALTH as this study lasted longer than NETSENSE. Fig. 1b shows the density of communication events with respect to the time of day when they were initiated. Most communication starts around noon and continues up to 11 PM. Fig. 1c plots the empirical cumulative distribution of the inter-event times – we observe that these inter-event times are long-tailed distributed, a result already known in literature [19]. Lastly, Fig. 1d and Fig. 1e present for NETSENSE and NETHEALTH, respectively, the number of communication peers for each of the students participating in the survey. We distinguish between incoming and outgoing peers and filter out the peers with less than 20 interactions to remove occasional phone calls. Visibly, for NETSENSE, the two density plots mostly overlap; however, for NETHEALTH, the students appear to have more incoming peers – probably reflecting the changes in how college students communicate as the conversations involved many short messages on platforms like iMessage. For both studies, most students have less than 130 peers, which is compatible with Dunbar’s number and the theory on cognitive limits of the human brain [10]. For NETSENSE,

out of the 3,159,669 total outgoing communication events, 2,012,816 have been interactions with peers that students labeled in surveys with the type of the relationship. That gives 1,779 relationships to investigate for NETSENSE. For NETHEALTH, out of 23,705,245 outgoing events, 15,605,131 were with 3,524 peers that have been categorized – the number of instances for each category is shown in Table 1. The following section fits a Hawkes model for each communication relationship.

D. ANALYZE CALL PATTERNS USING HAWKES

We model the call series between two individuals as a point process. We assume that future calls are more likely when other calls have recently occurred – in other words, the calling process has the self-exciting property. This hypothesis implies that the future of a sequence of events depends on its entire past and not only the current state. In other words, it does not have the Markov property. Consequently, typical choice modeling tools (such as Markov chains models) are unsuitable for modeling calls. Instead, we fit the call series using Hawkes point processes – which is a well-understood non-Markovian extension of the Poisson process [46]. We leverage a general-purpose point process R library (*evently* [24]) to produce user-level descriptions based solely on the fitted model parameters of the call series.

1) MAP CALLS TO POINT PROCESSES

In this work, we use phone calls and texts (both denoted hereafter as calls) as *point process events* (defined in [Section Prerequisites: Hawkes Processes](#)). Note that NETHEALTH includes the WhatsApp and iMessage group messages, which we remove since they are not peer-to-peer communication and are likely to follow other dynamics. An event in our datasets is a tuple (timestamp, sender, receiver). We build *event series* by grouping incoming and outgoing events between pairs of users. Each timestamp t_i is the time difference (in days) between the recorded timestamp of calls and the timestamp of the first call ($t_0 = 0$). We retain only the series where the sender is a student enrolled in the study. We impose the latter

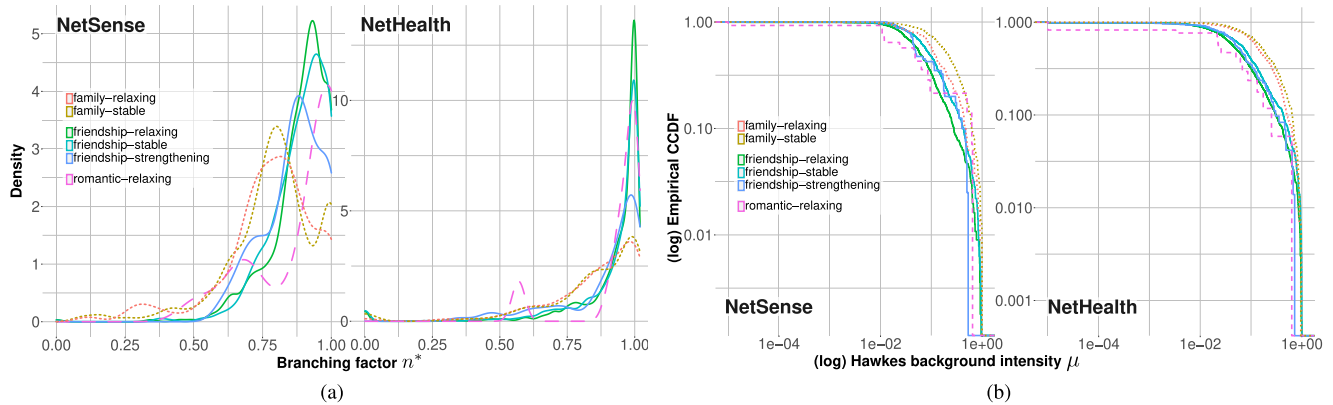


FIGURE 2. Distribution of Hawkes process parameters per relationship types: (a) the density plot for the branching factor n^* , (b) the complementary cumulative distribution function of the background event intensity parameter μ . Family relations are shown in dotted lines, friendship in solid lines, and romantic in dashed lines.

condition so that, in [Section Inferring Psychometric Traits](#), we can match the call series with the surveys filled in by the students. When both the sender and receiver are students in the study, the outgoing series for the sender is identical to the incoming series for the receiver. We filter our call series with fewer than 20 or more than 10,000 recorded calls or texts. We require the former condition to have enough data to fit the Hawkes processes (described next). The latter condition is to avoid computational explosion (considering that Hawkes is quadratic with the number of events). This results in 15,110 call series associated with 178 students, totaling more than 5 million call events for `NETSENSE`. For `NETHEALTH`, we get 70,752 call series associated with 712 students, totaling more than 33 million events.

2) FIT HAWKES PROCESSES

For each obtained call series, we fit the parameters of a Hawkes model, using the exponential and the power-law kernel functions ($\phi_{EXP}(t)$ and $\phi_{PL}(t)$), respectively, defined in [Section Prerequisites: Hawkes Processes](#)). We consider the exogenous intensity function to be constant ($\mu(t) = \mu$), and we fit it from data together with the other parameters using the software package `evently` [24]. `evently` is a R package for modeling events series using Hawkes processes and their variants. Internally, `evently` leverages `IPOPT` [59] – the current state of the art in constrained, non-linear optimization – to maximize the log-likelihood function in Eq (2). By design, it supports a wide array of kernel functions and provides an integrated set of functionalities to conduct event series-level or aggregated-level analyses. For each event series, `evently` outputs the fitted parameters μ , κ , θ (and c for $\phi_{PL}(t)$), and the branching factor n^* .

3) INTERPRETING THE HAWKES PARAMETERS

The Hawkes process is a generative model with interpretable parameters. The branching factor is a crucial quantity for point processes. It indicates the size of the self-exciting effect – when n^* is large, previous calls generate more future calls, and the recent calling history mainly drives the calling

TABLE 2. Percentage of call series (for each dataset) passing the goodness of fit tests (p -value > 0.05 , see details) and median holdout log-likelihood per event (LL) for the two Hawkes kernels. Higher is better. The winner kernel is shown in bold.

		KS	ED	LB	KS wins	LL
NETSENSE	Exponential	50.38%	37.30%	88.62%	20.90%	0.401
	Power-Law	68.12%	58.63%	82.49%	78.10%	0.665
NETHEALTH	Exponential	39.22%	31.03%	89.60%	22.50%	0.366
	Power-Law	61.74%	58.30%	87.59%	74.50%	0.666

relation. The background intensity indicates the size of the exogenous effect – when μ is large, calls are driven by outside factors. For example, one would expect that the call series between parents and children would exhibit a high μ , e.g., by calling daily at the same time. The parameter θ controls the decay rate of the call generation – when θ is small, the kernel function $\phi(t)$ decays slowly, and new calls are generated endogenously for more extended periods.

4) BUILD USER REPRESENTATION

We construct user descriptions using the methodology described by Kong et al. [26]. We jointly fit all call series involving a given user and describe each user’s activity with a latent mixture of Hawkes models. Due to Hawkes’ quadratic complexity, the joint fitting has significant computation speed advantages over a single large Hawkes model containing all user events. We build the user description by first computing the 5% percentiles bins for every parameter (i.e., μ , κ , c , and θ) for all mixture components and all users. Next, we compute the percentage of mixture parameters within each bin for each user. This results in 20 values (one per 5% percentile bin) for each parameter. In the end, we add a summary of temporal features for user series – the six-point summaries (min, mean, median, max, 25th and 75th percentile) of the call inter-arrival times, the total number of calls in a series and the duration of a series (in days). We also add the total number of call series in which a user is involved. This results in a 99-length vector describing user calling activity.

IV. SELECT A HAWKES PROCESS KERNEL

Here, we select the Hawkes kernel that best describes the time-decaying influence of calls in telecommunication series. We use goodness-of-fit tests and generalization error to compare the two kernels commonly used with human traces data – $\phi_{EXP}(t)$ the exponential and $\phi_{PL}(t)$ the power-law.

A. GOODNESS-OF-FIT TESTS

For point processes, the random time change theorem [3] states that the inter-event times transformed using the compensator (i.e., the definite integral of the event intensity) should follow a unit rate exponential distribution (see [26], [28] for more details). We perform three *portmanteau* statistical tests: the Kolmogorov-Smirnov (KS), Excess Dispersion (ED), and Ljung-Box (LB). For all three tests, to reject the null hypothesis (p-value < 0.05) means failing the test (that is, the sequence does not have the required property). The KS and ED tests assess whether the transformed interevent times are drawn from a unit exponential distribution, whereas LB tests their independence. We also compare the KS statistic for the exponential and power-law kernels to decide which fits better for the passing KS tests.

B. GENERALIZATION ERROR

We test the generalization of each kernel in a temporal holdout setup. We temporally split each call series into two parts. First, the earlier 80% of the call events in each series are used to train model parameters. Next, we compute the holdout log-likelihood on the later 20% of the events. Finally, we normalize the holdout log-likelihood by the number of events in the test set to account for the variable number of calls between series – i.e., we compute the log-likelihood per event.

C. RESULTS

Table 2 shows the percentage of call series that pass the statistical tests for each kernel and each dataset. Visibly, power-law exhibits a higher percentage of passing the fitness tests (KS and ED) than exponential. Both kernels show a high passing rate for the independence test (LB). Out of the passing KS tests, power-law wins the highest percentage of pairs when comparing the KS statistic (KS wins)—78.1% for $N_{ETSENSE}$ and 75.4% for $N_{ETHEALTH}$. Table 2 also shows the median holdout log-likelihood per event for each kernel and each dataset (we show it summarized as boxplots in the online appendix [1]). The power-law kernel outperforms the exponential kernel for both datasets. Together with the goodness-of-fit tests, this leads us to conclude that the power-law describes the call series better than exponential. Therefore, in the rest of this paper, we only present results for the power-law kernel with all call series refitted using all the available calls.

V. CHARACTERIZE RELATIONSHIPS

In *Section Describe relations using Hawkes*, we show that the fitted model parameters are descriptive for relationship types.

In *Section Classify relationship types*, we train classifiers to distinguish relationship types.

A. DESCRIBE RELATIONS USING HAWKES

Here, we use the Hawkes parameters' interpretability and analyze the fitted Hawkes parameters per relationship type. We concentrate on two quantities: the branching factor n^* (Fig. 2a) and the background event intensity μ (Fig. 2b). Fig. 2 shows that the different relations types have specific exogenous and self-excitation patterns. *Family relations* (shown in dotted lines) have lower branching factors and higher background intensity values. This indicates that outside factors mainly drive the calling dynamics between family members – e.g., students might call their families every day, at the same time, to let them know they are safe. *Friendship relations* (shown in solid lines) appear to have the opposite pattern. They are mainly driven by self-excitation (high n^*), and friends are less likely than family to call each other after a period of pause (low μ) – linking to the old adage: “out of sight, out of mind”. Finally, *romantic relations* (dashed lines) have a peculiar mixed behavior. Both n^* and μ show two modes, one for lower values and one for higher values, indicating a dual nature of romantic relations – some partners call each other because of recent calling activity; in contrast, others use more of a scheduled approach. These patterns are strongly consistent across both datasets and indicate that we can distinguish between types of relationships using a machine learning classifier (see next section).

B. CLASSIFY RELATIONSHIP TYPES

Here, we ask whether the relationships described using fitted Hawkes processes are identifiable one from another.

1) PREDICTIVE SETUP

We describe each relationship using the fitted Hawkes parameters μ , κ , θ and c , and its branching factor n^* . In addition, we train an ARIMA [31] model for each relationship and use the fitted parameters as a baseline to compare with the predictive performance of Hawkes features. Next, we use the Hawkes, ARIMA, and concatenated Hawkes+ARIMA feature sets to train three off-the-shelf classification algorithms: Random Forests [2], Support Vector Machines (SVM) [8], and XGBoost [5]. We compute the prediction performance of each algorithm using two nested cross-validation loops. The outer loop uses five folds to estimate the generalization performance on unseen data. The inner loop tunes hyperparameters using 5-fold cross-validation and random search with 500 combinations. We use the SMOTE algorithm to balance the size of the classes.

2) RESULTS

Fig. 3 shows the prediction performance for Random Forest, which achieved the best result (see the results for the other two classifiers in the online appendix [1]). The

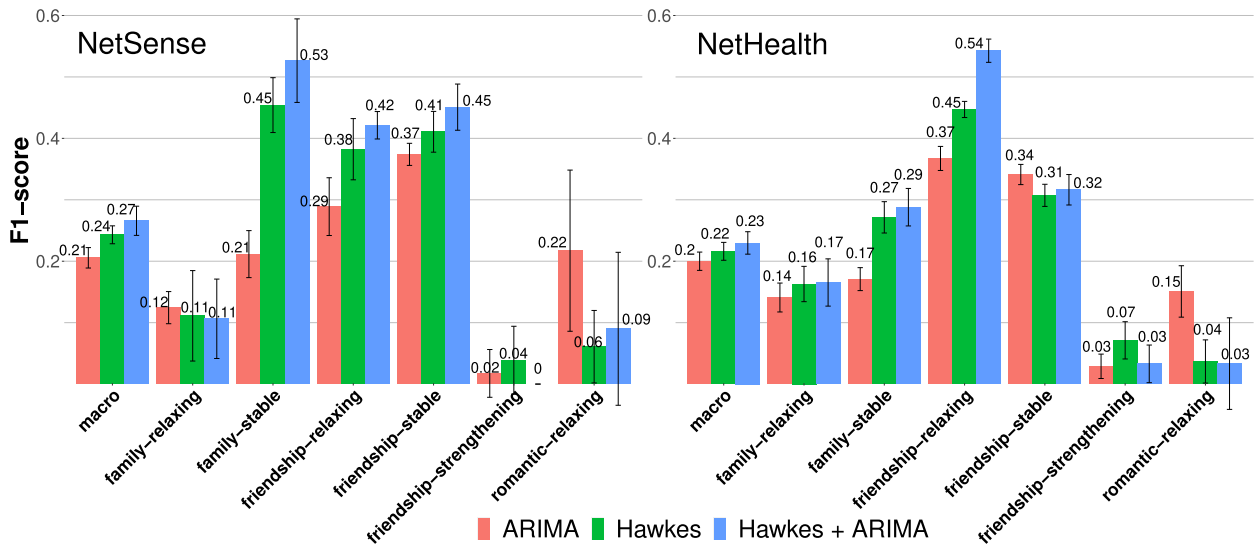


FIGURE 3. Performance of predicting relationship types using the call activity modeled using Hawkes processes and ARIMA (higher is better).

Hawkes features obtain a macro F1-score of 24% for $\text{NET}_{\text{SENSE}}$ and 22% for $\text{NET}_{\text{HEALTH}}$, more than double the random classifier and outperforming the ARIMA baseline (21% and 20%, respectively). Furthermore, the Hawkes and the ARIMA approaches are complementary, as the concatenated feature set Hawkes+ARIMA yields the best results. Unsurprisingly, the larger classes (*family-stable*, *friendship-stable* and *friendship-relaxing*) are consistently best predicted. It is perhaps surprising how poorly predicted is the *family-relaxing* class given its size, which indicates that there are multiple dynamics of how students interact increasingly less with their family. Overall, the results show that relationship types (and their dynamics) can be inferred from the parameter of a Hawkes process fitted on the call series. Additionally, we performed temporal change point detection in the dynamic relations using a paired Wilcoxon signed-rank test (conf. level = 95%) and effect size (Cohen's d [7]) to compare the log-likelihood before and after the change (more details in the online appendix [1]). We find that statistically significant results (p -value $\ll 0.001$) were obtained for $\text{NET}_{\text{HEALTH}}$ for *family-relaxing* (effect size = 0.25) and *friendship-relaxing* (effect size = 0.14).

VI. INFER PSYCHOMETRIC TRAITS

Here, we use call activity to infer the students' Big5 profiles [15]. First, in [Section Baselines and feature sets](#) we detail the baselines and sets of features. Next, in [Section Inferring psychometric traits](#), we present the prediction performance.

A. BASELINES AND FEATURE SETS

We set out to predict personality traits using three sets of features, which capture the different aspects of user activity. The first set is the Hawkes-based user representation that we introduced in [Section Analyzing call patterns using](#)

Hawkes. These features use the call series solely, and we denote them as *Hawkes*. The second set of features, denoted as *Monsted*, replicates the work of Mønsted et al. [38], who use autoregressive models to account for the historical dependency between current and historical volumes of calls in a series. They also build call and text statistics, such as duration, inter-event time, or volume. Unlike the original paper, we do not build features concerning online social network activity (Facebook), proximity (Bluetooth), and geospatial mobility (GPS) – for which the data is not fully available in our studies. The third set of features, denoted as *Surveys*, builds upon the surveys filled by students. For both $\text{NET}_{\text{SENSE}}$ and $\text{NET}_{\text{HEALTH}}$, we extract information that describes the students, such as their grades, perception of their health, number of books they read, or the number of clubs or social organizations they belong to. This results in 145 features for $\text{NET}_{\text{SENSE}}$ and 2212 features for $\text{NET}_{\text{HEALTH}}$. Note that *Surveys* use the information provided by the students, whereas *Hawkes* and *Monsted* use only externally observed information (i.e., call series). Note also that *Surveys* does not contain the forty-four questions related to the Big5 questionnaire (and which serve as ground truth for the predictive exercise). Finally, we also test combinations of feature sets to assess their complementarity.

B. INFERRING PSYCHOMETRIC TRAITS

1) PREDICTIVE SETUP

We generate the three sets of features for each of the 178 users for $\text{NET}_{\text{SENSE}}$ and 712 for $\text{NET}_{\text{HEALTH}}$ for whom the *Hawkes* features exist (see [Section Analyzing call patterns using Hawkes](#)). We trialed several off-the-shelf regressors – including SVM, XGBoost, KNN, and Random Forests – out of which Random Forests performed best and is shown in the rest of this section. We tune the hyperparameters and compute

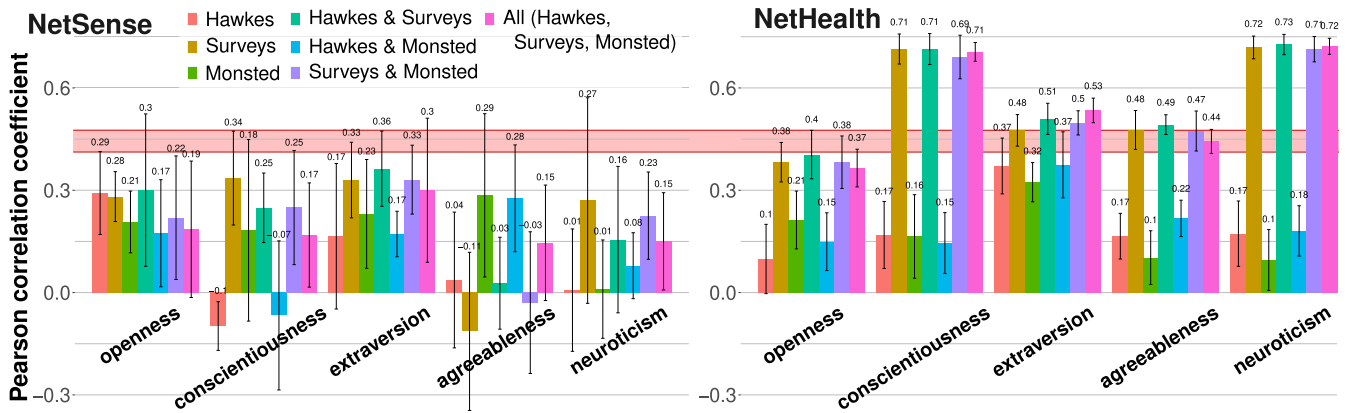


FIGURE 4. Big5 traits prediction performance measured using the Pearson correlation coefficient. Students are described using features built on two data sources two data sources (calls and filled-in surveys). For calls, we describe the users using our Hawkes-model embedding (Hawkes) and using a recent baseline [38] (Monsted). We also test combinations of feature sets. The red area shows the credible upper limits for correlations between predicted and self-reported personality traits [40].

prediction performance using 5-fold nested cross-validation and random search with 500 combinations.

2) PERFORMANCE OF *Hawkes* FEATURES

In line with recent literature [27], [38], [40], [52], we report in Fig. 4 personality traits prediction performances using Pearson’s correlation coefficient. We show the RMSE performances in the online appendix [1]. In their recent survey study, Novikov et al. [40] state that the upper limits for correlations between predicted and self-reported personality traits vary between 0.42 and 0.48. Our *Hawkes* features achieve 0.29 for *openness* on NET_{SENSE} , and 0.37 for *extraversion* on NET_{HEALTH} — very close to the upper limit range. The other traits are predicted around 0.17 on NET_{HEALTH} . For NET_{SENSE} , *Hawkes* performs poorly for *conscientiousness*, *agreeableness* and *neuroticism*.

3) NET_{SENSE} VS NET_{HEALTH}

We observe a higher variability of performance for NET_{SENSE} than for NET_{HEALTH} and a significantly higher variance. We posit this is due to the length of the studies – NET_{SENSE} it is only two and a half years, while for NET_{HEALTH} the period is four and a half years – and the amount of data available for each. Since personality traits do not change significantly over the years [40], longer communication periods may contribute to more accurate predictions. We conclude that NET_{HEALTH} is a better dataset for predicting personality traits from communication data.

4) SURVEYS FEATURES

Unsurprisingly, the best performing features overall are the *Surveys* which embed self-reported information – the gold standard in psychology. On NET_{SENSE} , they achieve around 0.3, and for NET_{HEALTH} it obtains more than 0.7 on *conscientiousness* and *neuroticism*. We make two observations. First, *Features* outperform the other feature sets given the richness of self-reported information (145 questions for

NET_{SENSE} and 2212 for NET_{HEALTH}). However, we note that surveying is a very costly process both financially and in time. Second, *Hawkes* and *Surveys* appear to capture complementary aspects of personality, as putting them together further increases performances, particularly for *openness* and *extraversion* on both datasets.

5) HAWKES VS MONSTED

For NET_{HEALTH} , the *Hawkes* features are outperforming the *Monsted* on four out of the five traits and even come close to the prediction based on the very complex *Surveys*. The most significant advantage of *Hawkes* over *Monsted* is that it does not require complex feature engineering, which may not be easily portable to new datasets.

VII. RELATED WORK

We structure related works discussion into two sections. First, we discuss work that used Hawkes processes to model contacts and interactions. Second, we explore works that infer personality traits from online data sources.

A. MODELING INTERACTIONS IN CONTACT NETWORKS USING POINT PROCESSES

Hawkes processes have been widely used to model social interactions in several applications because they can account for bursts of activity localized in time. Zipkin et al. [63] study electronic communications in a dataset of emails pertaining to the US military, where they observe that activity along the edges of the communication network is bursty. They apply a Hawkes model for the email exchange along edges and focus on studying parameter estimation in the presence of missing data. Moore and Davenport [39] learn the topology of a wireless network from limited passive observations of network activity. They use a multi-variate Hawkes process; they show it can detect changes to the existing topology and extract higher-level summaries of information flow in the network. Choudhari et al. [6] models simultaneously

events and the structure of a social network using a Hidden Markov Hawkes Process that incorporates topical Markov Chains within Hawkes processes to jointly model topical interactions along with the user-user and user-topic patterns. Hawkes processes have also been used to model face-to-face interactions in offices [35] as well as retweet cascades. For the latter task, Kobayashi and Lambiotte [23] propose a Time-Dependent Hawkes process to account for the circadian nature of the users and the aging of information when modeling retweet cascades, whereas Mishra et al. [37] leverages a Hawkes process with a power-law relaxation kernel and uses it with the user and timing features to predict the popularity of retweet cascades. Luo et al. [30] use company employees' phone calls to study the alignment method for Hawkes processes based on fused Gromov-Wasserstein discrepancy.

To the best of our knowledge, this is the first work that applies Hawkes modeling to telecom data and uses its outputs to predict users' relations and personality traits.

B. PREDICTING PERSONAL TRAITS FROM SOCIAL INTERACTIONS

A fertile area of computational psychology deals with inferring psychometric traits from a range of sources made available by our new interconnected society. For example, Settanni et al. [49] show that digital traces from social media can be studied to assess and predict theoretically distant psychosocial characteristics with remarkable accuracy. They also show that when additional user demographics are leveraged as additional digital traces, the accuracy of predictions improves. Some of these works also use social media-inferred personality traits to influence opinions and behavior or infer private traits. For example, Matz et al. [36] performed three field experiments that reached over 3.5 million individuals with psychologically tailored advertising and found that matching the content of persuasive appeals to individuals' psychological characteristics rendered the messaging significantly more effective. Kosinski et al. [27] used a public source of social media data (i.e., Facebook likes) to predict the Big5 personality traits of users, alongside a range of highly sensitive personal attributes, including sexual orientation, ethnicity, religious and political views, personality traits, intelligence, happiness, use of addictive substances, parental separation, age, and gender. A follow-up study by Youyou et al. [62] even showed that personality traits prediction based on Facebook likes are more accurate than the user's friends' estimations based on surveys.

The prior work most relevant to this study relates to learning personality traits from behavioral information collected via smartphones. Such works usually collect data from the onboard sensors and other phone logs [16]. The works closest to ours are by Stachl et al. [52], and Mønsted et al. [38]. Stachl et al. [52] predict Big Five personality dimensions using six different classes of behavioral information collected

from smartphones: 1) communication and social behavior, 2) music consumption, 3) app usage, 4) mobility, 5) overall phone activity, and 6) day- and night-time activity. They find that the accuracy of these predictions is similar to that of social media platforms. On the contrary, Mønsted et al. [38] claim that smartphone usage is not as predictive of Big5 personality traits as previously reported, and higher predictabilities in the literature are likely due to overfitting on small datasets.

Our work differs from those mentioned above in two major aspects. First, we use the Hawkes models fitted on the call series to predict user traits. As far as we are aware, no other work has used fitted Hawkes point processes to distill the call interactions between users and predict Big5 traits. Second, the above work requires access to the user's phone, as most features need to be recorded on the device. Our work shows that personality traits can be accurately predicted solely on call and text logs, which can be obtained outside the user's device. We also show that our Hawkes descriptors' prediction accuracy is comparable to that obtained from user-filled surveys.

VIII. CONCLUSION AND FUTURE WORK

This work investigates whether the Hawkes processes trained on telecommunication metadata predict human sociological traits and aspects. We employ two extensive datasets – NETSENSE and NETHEALTH, which contain mobile phone communication metadata and detailed information on surveyed university students. We fit a Hawkes process for each pair (participating student, peer) and use its fitted parameters as descriptors. In a series of experiments, we show that the Hawkes processes can distinguish between types of relationships, predict their temporal dynamics and infer user Big5 psychometric traits. This work is the first to show that Hawkes processes can be used to abstract detailed communication events that carry additional human sociological information. Due to the novelty of our application of the Hawkes process in the described problem, we decided that the right approach would be to start research with a simple model, the vanilla Hawkes process. Our work and results could provide a starting point for future research that will focus on using more complex point process models.

A. THE ETHICS OF USING TELECOM TRACES TO INFER PRIVATE TRAITS

These capabilities raise critical questions concerning user privacy. Our work joins a body of evidence [27], [50], [51] to show how the users' digital traits can be used to infer the user's sensitive information (such as the psychological profile). It can be disconcerting that such rich sources of private information lie outside the control of users and in the data warehouses of third parties, and we argue that its usage should be regulated. As Taylor [57] argues, studying human mobility using telecom data is a double-edged sword.

B. LIMITATIONS AND FUTURE WORK

For a given pair of individuals (sender, receiver), this work merges outgoing and incoming calls into a single call series. Future work could apply a bivariate Hawkes process to model the intertwining of both aspects of a discussion. In addition, as mentioned earlier, future work may focus on using more complex point process models. For example, call activity may be correlated with time, so it may be essential to use a model that takes into account the daily or weekly seasonality in the background rate [41]. Another direction of research is to use deep learning, however, this will result in a loss of interpretability of the model.

IX. DATA AVAILABILITY STATEMENT

The NETSENSE dataset is available upon request from Prof. Omar Lizardo, NETHEALTH data is publicly available for research purposes: <https://sites.nd.edu/nethealth/>. The code and the data samples are available online at <https://github.com/pwr-ai/predicting-relationships-and-big5-using-hawkes>.

REFERENCES

- [1] (2021). *Predicting Relationship Labels and Individual Personality Traits From Telecommunication History in Social Networks Using Hawkes Processes*. [Online]. Available: <https://bit.ly/3fWgsPP>
- [2] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.
- [3] E. N. Brown, R. Barbieri, V. Ventura, R. E. Kass, and L. M. Frank, "The time-rescaling theorem and its application to neural spike train data analysis," *Neural Comput.*, vol. 14, no. 2, pp. 325–346, 2002.
- [4] R. S. Burt, "Decay functions," *Social Netw.*, vol. 22, no. 1, pp. 1–28, May 2000.
- [5] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. KDD*, Aug. 2016, pp. 785–794.
- [6] J. Choudhari, A. Dasgupta, I. Bhattacharya, and S. Bedathur, "Discovering topical interactions in text-based cascades using hidden Markov Hawkes processes," in *Proc. ICDM*, Nov. 2018, pp. 923–928.
- [7] J. Cohn, *Statistical Power Analysis for the Behavioral Sciences*. Hillsdale, NJ, USA: Lawrence Erlbaum Associates, 1988.
- [8] C. Cortes and V. Vapnik, "Support-vector networks," *Mach. Learn.*, vol. 20, no. 3, pp. 273–297, 1995.
- [9] D. J. Daley and D. Vere-Jones, "Conditional intensities and likelihoods," in *An Introduction to the Theory of Point Processes*, vol. 1. New York, NY, USA: Springer, 2008, ch. 7.2.
- [10] R. I. M. Dunbar, "Neocortex size as a constraint on group size in primates," *J. Human Evol.*, vol. 22, no. 6, pp. 469–493, Jun. 1992.
- [11] J. Etesami, W. Trouleau, N. Kiyavash, M. Grossglauser, and P. Thiran, "A variational inference approach to learning multivariate wold processes," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2021, pp. 2044–2052.
- [12] S. L. Feld, J. J. Sutor, and J. G. Hoegh, "Describing changes in personal networks over time," *Field Methods*, vol. 19, no. 2, pp. 218–236, May 2007.
- [13] J. Flamino, R. DeVito, B. K. Szymanski, and O. Lizardo, "A machine learning approach to predicting continuous tie strengths," 2021, *arXiv:2101.09417*.
- [14] E. Gilbert and K. Karahalios, "Predicting tie strength with social media," in *Proc. SIGCHI Conf. Hum. Factors Comput. Syst.*, Apr. 2009, pp. 211–220.
- [15] L. R. Goldberg, "An alternative 'description of personality': The big-five factor structure," *J. Pers. Social Psychol.*, vol. 59, no. 6, pp. 1216–1229, 1990.
- [16] G. M. Harari, R. S. Müller, C. Stachl, R. Wang, W. Wang, M. Bühner, P. J. Rentfrow, A. T. Campbell, and S. D. Gosling, "Sensing sociability: Individual differences in young adults' conversation, calling, texting, and app use behaviors in daily life," *J. Pers. Social Psychol.*, vol. 119, no. 1, pp. 204–228, 2019.
- [17] A. G. Hawkes, "Spectra of some self-exciting and mutually exciting point processes," *Biometrika*, vol. 58, no. 1, pp. 83–90, 1971.
- [18] G. Alan Hawkes and D. Oakes, "A cluster process representation of a self-exciting process," *J. Appl. Probab.*, vol. 11, no. 3, p. 493, Sep. 1974.
- [19] H.-H. Jo, M. Karsai, J. Kertész, and K. Kaski, "Circadian pattern and burstiness in mobile phone communication," *New J. Phys.*, vol. 14, no. 1, Jan. 2012, Art. no. 013055.
- [20] O. P. John, L. P. Naumann, and C. J. Soto, "Paradigm shift to the integrative big five trait taxonomy: History, measurement, and conceptual issues," in *Handbook of Personality: Theory and Research*, 3rd ed. New York, NY, USA: Guilford Press, 2008, pp. 114–158.
- [21] D. H. Johnson, "Point process models of single-neuron discharges," *J. Comput. Neurosci.*, vol. 3, no. 4, pp. 275–299, 1996.
- [22] M. L. Kern, P. X. McCarthy, D. Chakrabarty, and M.-A. Rizoïu, "Social media-predicted personality traits and values can help match people to their ideal jobs," *Proc. Nat. Acad. Sci. USA*, vol. 116, no. 52, pp. 26459–26464, Dec. 2019.
- [23] R. Kobayashi and R. Lambiotte, "TiDeH: Time-dependent Hawkes process for predicting retweet dynamics," in *Proc. ICWSM*, 2016, pp. 1–10.
- [24] Q. Kong, R. Ram, and M.-A. Rizoïu, "Evently: Modeling and analyzing reshare cascades with Hawkes processes," in *Proc. WSDM*, Mar. 2021, pp. 1097–1100.
- [25] Q. Kong, M.-A. Rizoïu, and L. Xie, "Describing and predicting online items with reshare cascades via dual mixture self-exciting processes," in *Proc. CIKM*, Oct. 2020, pp. 645–654.
- [26] Q. Kong, M.-A. Rizoïu, and L. Xie, "Modeling information cascades with self-exciting processes via generalized epidemic models," in *Proc. WSDM*, Jan. 2020, pp. 286–294.
- [27] M. Kosinski, D. Stillwell, and T. Graepel, "Private traits and attributes are predictable from digital records of human behavior," *Proc. Nat. Acad. Sci. USA*, vol. 110, no. 15, pp. 5802–5805, Apr. 2013.
- [28] P. J. Laub, T. Taimre, and P. K. Pollett, "Hawkes processes," 2015, *arXiv:1507.02822*.
- [29] S. Liu, D. Hachen, O. Lizardo, C. Poellabauer, A. Striegel, and T. Milenković, "Network analysis of the NetHealth data: Exploring co-evolution of individuals' social network positions and physical activities," *Appl. Netw. Sci.*, vol. 3, no. 1, pp. 1–26, Dec. 2018.
- [30] D. Luo, H. Xu, and L. Carin, "Fused Gromov–Wasserstein alignment for Hawkes processes," 2019, *arXiv:1910.02096*.
- [31] S. Makridakis and M. Hibon, "ARMA models and the Box–Jenkins methodology," *J. Forecasting*, vol. 16, no. 3, pp. 147–163, May 1997.
- [32] R. D. Malmgren, D. B. Stouffer, A. E. Motter, and L. A. N. Amaral, "A Poissonian explanation for heavy tails in e-mail communication," *Proc. Nat. Acad. Sci. USA*, vol. 105, no. 47, pp. 18153–18158, Nov. 2008.
- [33] H. Mao, X. Shuai, and A. Kapadia, "Loose tweets: An analysis of privacy leaks on Twitter," in *Proc. ACM CCS*, Oct. 2011, pp. 1–11.
- [34] B. Marcus, F. Machilek, and A. Schütz, "Personality in cyberspace: Personal web sites as media for personality expressions and impressions," *J. Pers. Social Psychol.*, vol. 90, no. 6, pp. 1014–1031, Jun. 2006.
- [35] N. Masuda, T. Takaguchi, N. Sato, and K. Yano, "Self-exciting point process modeling of conversation event sequences," in *Temporal Networks*. Berlin, Germany: Springer, 2013.
- [36] S. C. Matz, M. Kosinski, G. Nave, and D. J. Stillwell, "Psychological targeting as an effective approach to digital mass persuasion," *Proc. Nat. Acad. Sci. USA*, vol. 114, no. 48, pp. 12714–12719, Nov. 2017.
- [37] S. Mishra, M.-A. Rizoïu, and L. Xie, "Feature driven and point process approaches for popularity prediction," in *Proc. CIKM*, Oct. 2016, pp. 1069–1078.
- [38] B. Mønsted, A. Mollgaard, and J. Mathiesen, "Phone-based metric as a predictor for basic personality traits," *J. Res. Pers.*, vol. 74, pp. 16–22, Jun. 2018.
- [39] M. G. Moore and M. A. Davenport, "Analysis of wireless networks using Hawkes processes," in *Proc. IEEE 17th Int. Workshop Signal Proc. Adv. Wireless Commun.*, Jul. 2016, pp. 1–5.
- [40] P. Novikov, L. Mararitsa, and V. Nozdachev, "Inferred vs traditional personality assessment: Are we predicting the same thing? Mar. 2021, *arXiv:2103.09632*.
- [41] T. Omi, Y. Hirata, and K. Aihara, "Hawkes process model with a time-dependent background rate and its application to high-frequency financial data," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 96, no. 1, Jul. 2017, Art. no. 012303.

- [42] D. J. Ozer and V. Benet-Martínez, "Personality and the prediction of consequential outcomes," *Annu. Rev. Psychol.*, vol. 57, no. 1, pp. 401–421, Jan. 2006.
- [43] R. Purta, S. Mattingly, L. Song, O. Lizardo, D. Hachen, C. Poellabauer, and A. Striegel, "Experiences measuring sleep and physical activity patterns across a large college cohort with fitbits," in *Proc. ACM Int. Symp. Wearable Comput.*, Sep. 2016, pp. 28–35.
- [44] T. Raeder, O. Lizardo, D. Hachen, and N. V. Chawla, "Predictors of short-term decay of cell phone contacts in a large scale communication network," *Social Netw.*, vol. 33, no. 4, pp. 245–257, Oct. 2011.
- [45] M. T. Rivera, S. B. Soderstrom, and B. Uzzi, "Dynamics of dyads in social networks: Assortative, relational, and proximity mechanisms," *Annu. Rev. Sociol.*, vol. 36, no. 1, pp. 91–115, Jun. 2010.
- [46] M.-A. Rizoïu, A. Soen, S. Li, P. Calderon, L. Dong, A. K. Menon, and L. Xie, "Interval-censored Hawkes processes," *J. Mach. Learn. Res.*, vpl. 23, pp. 1–84, Apr. 2022.
- [47] M.-A. Rizoïu, L. Xie, T. Caetano, and M. Cebrian, "Evolution of privacy loss in Wikipedia," in *Proc. WSDM*, Feb. 2016, pp. 215–224.
- [48] B. W. Roberts, N. R. Kuncel, R. Shiner, A. Caspi, and L. R. Goldberg, "The power of personality: The comparative validity of personality traits, socioeconomic status, and cognitive ability for predicting important life outcomes," *Perspect. Psychol. Sci.*, vol. 2, no. 4, pp. 313–345, Dec. 2007.
- [49] M. Settanni, D. Azucar, and D. Marengo, "Predicting individual characteristics from digital traces on social media: A meta-analysis," *Cyberpsychol., Behav., Social Netw.*, vol. 21, no. 4, pp. 217–228, Apr. 2018.
- [50] C. Smith-Clarke, A. Mashhadi, and L. Capra, "Poverty on the cheap: Estimating poverty maps using aggregated mobile communication networks," in *Proc. CHI*. New York, NY, USA: Association for Computing Machinery, Apr. 2014, pp. 511–520.
- [51] V. Soto, V. Frias-Martinez, J. Virseda, and E. Frias-Martinez, "Prediction of socioeconomic levels using cell phone records," in *Proc. UMAP*, 2011, pp. 377–388.
- [52] C. Stachl, Q. Au, R. Schoedel, S. D. Gosling, G. M. Harari, D. Buschek, S. T. Völkel, T. Schuwerk, M. Oldemeier, T. Ullmann, H. Hussmann, B. Bischl, and M. Bühner, "Predicting personality from patterns of behavior collected with smartphones," *Proc. Nat. Acad. Sci. USA*, vol. 117, no. 30, pp. 17680–17687, Jul. 2020.
- [53] A. Striegel, S. Liu, L. Meng, C. Poellabauer, D. Hachen, and O. Lizardo, "Lessons learned from the netsense smartphone study," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 4, pp. 51–56, Sep. 2013.
- [54] J. J. Suiitor, B. Wellman, and D. L. Morgan, "It's about time: How, why, and when networks change," *Social Netw.*, vol. 19, no. 1, pp. 1–7, Jan. 1997.
- [55] J. Suiitor and S. Keeton, "Once a friend, always a friend? Effects of homophily on women's support networks across a decade," *Social Netw.*, vol. 19, no. 1, pp. 51–62, Jan. 1997.
- [56] T. Takaguchi, N. Sato, K. Yano, and N. Masuda, "Importance of individual events in temporal networks," *New J. Phys.*, vol. 14, no. 9, Sep. 2012, Art. no. 093003.
- [57] L. Taylor, "No place to hide? The ethics and analytics of tracking mobility using mobile phone data," *Environ. Planning D, Soc. Space*, vol. 34, no. 2, pp. 319–336, Apr. 2016.
- [58] J. Ureña-Carrion, J. Saramäki, and M. Kivelä, "Estimating tie strength in social networks using temporal communication data," *EPJ Data Sci.*, vol. 9, no. 1, p. 37, Dec. 2020.
- [59] A. Wächter and L. T. Biegler, "On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming," *Math. Program.*, vol. 106, no. 1, pp. 25–57, 2006.
- [60] B. Wellman, R. Y.-L. Wong, D. Tindall, and N. Nazer, "A decade of network change: Turnover, persistence and stability in personal communities," *Social Netw.*, vol. 19, no. 1, pp. 27–50, Jan. 1997.
- [61] T. Yarkoni, "Personality in 100,000 words: A large-scale analysis of personality and word use among bloggers," *J. Res. Pers.*, vol. 44, no. 3, pp. 363–373, Jun. 2010.
- [62] W. Youyou, M. Kosinski, and D. Stillwell, "Computer-based personality judgments are more accurate than those made by humans," *Proc. Nat. Acad. Sci. USA*, vol. 112, no. 4, pp. 1036–1040, Jan. 2015.
- [63] J. R. Zipkin, F. P. Schoenberg, K. Coronges, and A. L. Bertozzi, "Point-process models of social network interactions: Parameter estimation and missing data recovery," *Eur. J. Appl. Math.*, vol. 27, no. 3, pp. 502–529, Jun. 2016.



MATEUSZ NUREK is currently pursuing the Ph.D. degree with the Department of Artificial Intelligence, Wrocław University of Science and Technology, Wrocław, Poland, and a member of the Network Science Laboratory. His research interests include network science and machine learning. He focuses on applying data science techniques in the field of human behavior and social network analysis.



RADOŚLAW MICHALSKI (Member, IEEE) is currently an Associate Professor with the Department of Artificial Intelligence, Wrocław University of Science and Technology, Wrocław, Poland. He co-leads the Network Science Laboratory, Wrocław University of Science and Technology, and leads the Blockchain Exploration Research Group (BERG). His research interests include social influence, diffusion processes in complex networks, and machine learning. He has coauthored more than 50 publications in these areas.



OMAR LIZARDO is currently a Professor and the LeRoy Neiman Term Chair with the University of California at Los Angeles, Los Angeles, CA, USA. He studies culture, cognition, networks, consumption, institutions, organization, and social theory. He serves on the Editorial Board of the social science journals, such as *Social Forces*, *Sociological Theory*, *Sociological Forum*, *Journal for the Theory of Social Behaviour*, *Theory and Society*, and *Poetics*. He is also a member of the Board of Reviewing Editor for the journal *Science* and an Associate Editor of *Discover Data*.



MARIAN-ANDREI RIZOIU is currently an Assistant Professor with the University of Technology Sydney, where he leads the Behavioral Data Science Group. He is also a Computer Scientist interested in addressing research questions where human online behavior and machine learning cross over. His research has made several key contributions to online popularity prediction, real-time tracking and countering disinformation campaigns, and understanding shortages and mismatches in labor markets. His research is funded by Facebook Research and the Commonwealth of Australia, and published in the most selective venues, such as the *PNAS*, *PLOS ONE*, *PLOS Computational Biology*, *WWW*, *NeurIPS*, *IJCAI*, and *CIKM*. In addition, his work has received significant media attention, including *Bloomberg Business Week*, *Nature Index*, *BBC*, and *World Economic Forum*. His work has a societal impact for social good, he serves as an expert in legislative initiatives and parliamentary inquiries.

• • •