

One Size Does Not Fit All: Profiling Personalized Time-Evolving User Behaviors

Pravallika Devineni¹, Evangelos E. Papalexakis¹, Danai Koutra², A. Seza Dođruöz³ and Michalis Faloutsos¹

¹Department of Computer Science & Engineering, University of California, Riverside, CA

²Department of Computer Science & Engineering, University of Michigan, Ann Arbor, MI

³Independent Researcher

pdevi002@ucr.edu

{epapalex, michalis}@cs.ucr.edu

dkoutra@umich.edu

a.s.dogruoz@gmail.com

Abstract—Given the set of social interactions of a user, how can we detect changes in interaction patterns over time? While most previous work has focused on studying network-wide properties and spotting outlier users, the dynamics of individual user interactions remain largely unexplored. This work sets out to explore those dynamics in a way that is minimally invasive to privacy, thus, avoids to rely on the textual content of user posts—except for validation. Our contributions are two-fold. First, in contrast to previous studies, we challenge the use of a fixed interval of observation. We introduce and empirically validate the “Temporal Asymmetry Hypothesis”, which states that appropriate observation intervals should vary both among users and over time for the same user. We validate this hypothesis using eight different datasets, including email, messaging, and social networks data. Second, we propose iNET, a comprehensive analytic and visualization framework which provides personalized insights into user behavior and operates in a streaming fashion. iNET learns personalized baseline behaviors of users and uses them to identify events that signify changes in user behavior. We evaluate the effectiveness of iNET by analyzing more than half a million interactions from Facebook users. Labeling of the identified changes in user behavior showed that iNET is able to capture a wide spectrum of exogenous and endogenous events, while the baselines are less diverse in nature and capture only 66% of that spectrum. Furthermore, iNET exhibited the highest precision (95%) compared to all competing approaches.

I. INTRODUCTION

User behavior modeling in social networks mostly focuses on the network-wide properties of the group behavior within a fixed time interval and aims at spotting users who are outliers [1], [2]. Some studies have focused on identifying network-wide or other static properties of users, such as graph patterns in the friendship graph [3], [4]. However, there is also need to focus on individual users who may deviate from their usual patterns of behaviors in social media. Few studies have used social media to identify postpartum depression in new mothers [5] or how quitting an addictive habit (e.g., smoking) is reflected on the social media behavior [6]. We discuss previous work in more detail in Section VI.

In this paper, we focus on the automatic and personalized identification of transitions in online user behaviors without accessing textual information and when the information arrives in a streaming fashion. The identified changes in social media behavior may point to real events and changes, some of which can benefit from intervention. Although exploring the

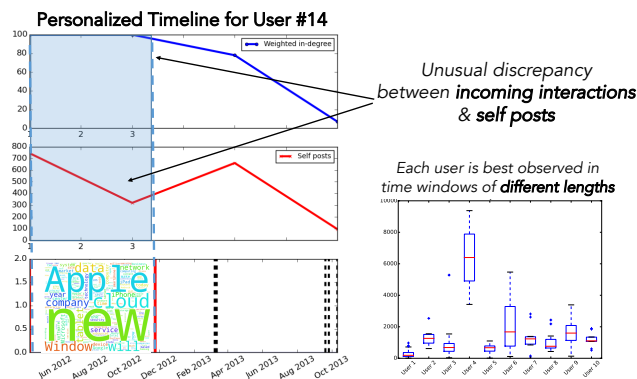


Fig. 1. iNET detects variable-duration periods of activity that deviates from the expected, across multiple streaming features of the user interaction (showing for Left: top: weighted in-degree, middle: self posts, bottom: keywords). In this particular example, the time period where the user’s behavior deviates from the expected pattern appears to be related to a product release or announcement by Apple. Right: Empirical validation of Temporal Asymmetry Hypothesis that states - different users have different time granularity and it varies over time.

correlation between online user behavior and real life events is an important next step, it is beyond the scope of this paper.

So far, most research involving time-evolving graphs consider *fixed intervals* of observation (e.g., a day or week) when they build the user interaction graphs. However, selecting the right interval of observation is challenging since different intervals may lead to different inferences from the data, observations and interventions. Instead, we propose a streaming method to detect changes in the online behaviors of the users with a user-centric approach.

Our work makes two main contributions. First, we challenge the widely adopted fixed interval of observation in favor of a variable interval of observation. An obvious challenge is how to adaptively determine these intervals. Second, we develop iNET, an analytic framework, to detect and visualize changes in online behaviors of individual users without using their text based interactions.

We use eight different real datasets to validate our ideas and techniques. These datasets include social media, email, and online forums and span multiple years of observation and up to more than a million interactions. We describe our datasets

in more detail in section II.

For illustration purposes, in Figure 1 we show an indicative result from using iNET on our Facebook dataset: The depicted user exhibits one window that deviates from the expected behavior (highlighted in blue) which, after further investigation for validation purposes, appears to be related to a new product release or announcement by Apple.

Our work can be summarized in the following points.

- **Temporal Asymmetry Hypothesis:** We introduce and empirically validate two hypotheses which state that the appropriate observation interval should vary among users and over time for the same user. This is a fundamental observation that forms the basis for the remainder of our work.
- **Mining & Visualization Framework:** We propose iNET, an algorithm that operates on streams of user interactions and detects periods of time for which there was an unusual deviation in user’s nature of interaction. iNET introduces a systematic way to identify variable intervals of observation in an personalized and adaptive way. We also present a comprehensive visual summary of the user’s timeline, highlighting the periods of interest, enabling an analyst to investigate the causes of the behavior further.
- **Representation Learning:** iNET learns a representation of a user’s social behavior over time by passing the limitations of the Temporal Asymmetry Hypothesis, since it uses variable length intervals of observation.
- **Case Study:** We evaluate the effectiveness of iNET by analyzing data from Facebook users. Our method was able to capture 100% of relevant categories, whereas baselines ranged between 22% to 66%. Furthermore, iNET exhibited the highest precision (95%) compared to all baselines.

Our work in perspective. We view our work as first step towards fully exploring the amount and type of information that can be extracted from the online social footprint of a person. We see two next steps: (a) identifying the minimum information needed to detect behavioral change, and (b) mapping online behavioral changes to real-life changes of the user. The first step is our effort to collect only the information needed, while the second step could have significant societal impact in detecting depression, manic episodes, suicidal tendencies, and cyber-bullying.

In the remainder of this paper, we first discuss the problem definition, the motivation and reasoning behind our work in more detail. Then, we describe the iNET algorithmic framework, our experiments and results. Finally, we present some related work and our conclusions.

II. PROBLEM DEFINITION AND DATASETS

We first present some terms and definitions and describe the datasets that we will use in our work.

Problem 1: Given a stream of interactions for a set of users, identify periods of time for which there was an significant deviation in a user’s nature of interaction from the expected behavior.

We represent the time-evolving interactions between users with “dynamic interaction networks”, i.e., a time series of graphs, each of which represents a graph snapshot: an aggregation of streaming interactions over a time interval of observation.

Graph snapshot at interval T . Given a time interval T , we represent the user interactions during T with a weighted directed graph $G_T(V_T, E_T)$. The edge direction indicates the “sender-receiver” flow of the interaction. The weight of an edge encodes the number of interactions between the nodes. Given the streaming setup of our problem, the interactions arrive at different times, and an edge is considered if it happens during the interval T . Similarly, a node is included in G_T , if it participates in an interaction during that interval.

As we mentioned earlier, identifying the right interval of observation is a challenge and we discuss it below.

User centric monitoring. Here, we are interested in understanding personalized interaction networks of users. As a result, we can create user-centric graphs (often called ego-nets), which are star graphs with the user of interest in its center (hub).

Example: For illustration purposes, we use a Facebook user interaction network. Every Facebook user has a wall, which can essentially be described as the user’s space. The wall captures all the ideas and thoughts shared by the user, and the user’s friends interact through likes, comments and posts on user’s wall. These interactions are time-labeled and directed and we consider them as *incoming interactions* since they are directed to the user. If the user authors a post on their own wall, we call this a *self post*. Given a start time and an end time, we can get a count of the number of self posts on a user’s wall in that period.

Our datasets. We describe the 8 datasets that we use in our work. Each dataset is a collection of users and the temporal interactions between them at a given timestamp. We consider all edges that involve a user’s participation to create the interaction network of the user.

- **Facebook1 (FB1) and Facebook2 (FB2), D2y:** We consider all interactions (e.g., comments, likes and wall-to-wall posts between a user and user’s friends) on Facebook occurring between user and user’s friends. FB1 and FB2 are of four-month periods in 2011 and 2012, respectively. The duration of D2y is 18.5 months between April 2012 and October 2013.
- **Enron:** This is an email interaction network of users [7] in an organization.
- **FB MPI:** This is an activity network of Facebook users writing on each others walls in 2009.
- **Digg:** This is the reply network of the social news website Digg. Each node is a user of the website, and each directed edge denotes that a user replied to another user.
- **Slashdot:** This is the reply network of technology website Slashdot. Nodes are users and edges are replies. The edges are directed and start from the responding user. Edges are annotated with the timestamp of the reply.

- **UC Irvine messages:** This directed network contains sent messages between the users of an online community of students from the University of California, Irvine. A node represents a user, and a directed edge corresponds to a sent message.

We accessed the FB MPI, Digg, Slashdot, and UC Irvine messages network datasets from the Koblenz network dataset repository [8]. The data in FB1, FB2, and D2y datasets was collected between 2011 and 2013 *via* a Facebook third-party application.

III. THE TEMPORAL ASYMMETRY HYPOTHESIS

The overarching goal of this paper is to observe change-points in a user’s behavior. Given the variability in a user’s behavior, we hypothesize that every user’s temporal behavior needs to be observed at a different time granularity, and further this granularity varies over time.

Several previous works used fixed-length intervals to analyze temporal data. Researchers in [9] analyzed the evolution of Facebook activity network by partitioning interactions between users into per day and per month intervals. To spot anomalies in network traffic measurements, [10] aggregated the data into 1-minute intervals. [11] used moving windows of size 7 days to detect deviations in user behavior in a large mobile phone network. Admittedly, some of our own previous works have used intervals of fixed lengths to analyze temporal data [2]. The amount of structure one requires in a network depends on what one intends to do with that network. Finding change in nature of interactions in a user’s behavior requires developing an insight into the underlying structure of interactions in time. Since every user is unique in terms of temporal behavior, each user’s interaction data stream has different rates. Aggregation lengths are often somewhat arbitrary, usually of a fixed-length in time, chosen based on intuition, convenience, or convention.

The limitations of fixed observations intervals. Choosing network aggregation intervals based on intuition has three short-comings. First, if the interval is too short, meaningful connections in time might be lost and the graph may lack sufficient structure for analysis. Second, if the interval is too long, we might remove the fine-grained changes critical to understanding temporal changes in user’s behavior. Third, if the user’s interaction exhibit a variable rates, fixed-length intervals will only be appropriate for part of the time of the experiment. In addition, if we analyze interactions over the entire duration of the dataset, we remove a lot of noise, smoothing out considerable amount of temporal information useful towards understanding the evolving nature of the user.

We hypothesize that fixed-length intervals do not provide enough support towards understanding time-evolving behavior of users.

Hypothesis 1: The appropriate time granularity of observation is dependent on the streaming rate of user interactions and varies from user to user.

To examine the validity of hypothesis 1, we conducted experiments using eight longitudinal network datasets from

a variety of domains. Each dataset consists of a set of users and timestamped interactions between them. Researchers have proposed frameworks [12], [13] to identify appropriate aggregation lengths in temporal streams. However powerful [12] is, it requires the entire set of interactions up front and, thus, we cannot use it in a streaming fashion.

[14] introduce ADAGE, which takes as input streaming graph edges of user interactions and partitions a user’s timeline into disjoint, variable-length intervals. It uses a network measure as part of its stopping criterion to detect the convergence of timeline into intervals. When we are dealing with a global graph (as in the case of [14]), there is a wide variety of network measures that we can employ in order to measure stability. However, in our case of a personalized interaction network, the graph we observe is essentially a *star* graph (i.e., the user of interest node is in the center and there are incoming and outgoing edges to and from that node). In this case, the number of possible metrics that we can use is heavily reduced by the type of graph we observe, and in fact, one of the most sensible and intuitive choices is the weighted degree, which captures the density of interaction. In most of the cases we studied, we only observed incoming interaction, and thus, we use the *weighted in-degree* as our metric of choice throughout the case study of this paper.

Formally, the problem that we would like to solve in order to identify the windows of observation per user is the following:

Problem 2 (Constructing graphs from streaming interactions): **Given** a stream of time-labeled edges for user interactions **construct** a snapshot of the graph for which a particular metric of the graph is stable.

There are two main aspects to Problem 2: Firstly, the temporal edges occur in a streaming fashion. Secondly, there is need to determine when a set of aggregated edges converges into a stable graph. Fortunately, [14] has equipped us with a methodology for handling the above two aspects and solving Problem 2.

Using ADAGE, we computed the appropriate intervals of individual users for the above datasets. Figure 2 presents the median of interval sizes for all users in a dataset computed based on the interaction streams of each user. The diversity




Fig. 2. Median of interval lengths of all users in the datasets measured in terms of weighted in-degree. High variability indicates asymmetry in temporal behavior of users.




Fig. 3. Variable-length intervals calculated for ten random users in various datasets. The distribution of quartiles indicates the high variability in the granularity of intervals within the users. The mere fact of this variability demonstrates Hypothesis 2.

in median values demonstrates that every user needs to be observed at a different time granularity, hence indicating that Hypothesis 1 is true. We hypothesize that even within an individual user, her behavior changes in time:

Hypothesis 2: The appropriate granularity of observation for a user varies over time.

Figure 3 illustrates the distribution of variable-length intervals for ten random users in the dataset. This implies that the rate of change of a user’s timeline varies over time. Hence, to understand emergence of behavior in individual users, variable-length intervals are required. In the next Section, we present our proposed method, which takes into account Temporal Asymmetry Hypothesis, learns a personalized representation for each user and conducts anomaly detection on that representation.

IV. PROPOSED METHOD

As we saw in the previous section, our overarching goal outlined by Problem 1 is to identify periods of time for which there was a significant deviation of a user’s nature of interaction from what is considered expected interaction behavior. Before we proceed with the description of the inner workings of our proposed method iNET, we introduce an instance of the problem as a case study. This case study will guide the method description. In order to do so, in the following lines we define expected and outlying interaction patterns for this particular instance of the problem.

Case Study: As we mention in the introduction, we focus on interaction-based networks that support functionalities of broadcasting a message to one’s peers and receiving interactions from those peers. Following the Facebook example, broadcasting is when a user posts on their own timeline,

whereas incoming interaction is anything that this user’s friend posts either on their timeline or pertaining to a particular “self-post”. For those types of networks that we study, we define expected and outlying interaction patterns as follows:

Definition 1 (Expected Interaction Pattern): The number of self-posts and the number of incoming interactions from the user’s peers are correlated.

Definition 2 (Outlying Interaction Pattern): An interaction pattern is defined as outlying if either (1) the number of self-posts is much higher than the number of incoming interactions, or, conversely, (2) the incoming interactions significantly outnumber the self-posts.

Outlying interactions, and changes in them, may indicate a variety of events, ranging from life changes to psychological disorders. Thus, being able to detect them early can be critical for timely intervention (especially in the latter case). We must note here that our definition of expected and outlying interaction patterns defines the particular choices we make in the metrics that iNET will use, however, as we stress earlier, iNET is modular and one can easily define a different “anomaly signature” and incorporate that to iNET.

A. Method Description

In the following lines we present iNET in detail, outlining and justifying particular design choices on the way. An outline of iNET is presented in Algorithm 1. iNET consists of the following three steps. Within each step, we first provide a general description that is agnostic to our particular case study, and subsequently we adapt each step to our case study.

Step 1: Representation Learning. For each user, we learn a personalized representation of their interaction patterns over time. As we saw in the previous section, due to the Temporal Asymmetry Hypothesis, one size for window of observation does not fit all users, therefore, we treat each user individually and we iteratively solve Problem 2 as in [14] in order to obtain a set of windows of observation for that user. After we have exhausted the entire stream of interactions for a particular user, we end up with a set of (w_i, f_i) values where w_i is the window of observation and f_i is the value of the network measure at the end of that window. Note that we can run this step for each user in parallel, or even interleave the representation learning for different users as we observe the stream of edges for all the users.

Case Study: The stability measure f_i of our graph, in this particular case, is the *weighted in-degree*. In order to be able to detect anomalous changes in the nature of interaction, according to Definitions 1 and 2, we also have to record the number of self-posts or broadcasts for a given user at the end of each window w_i . This is a simple counter that we need to keep track of in addition to the network measure f_i and at the end of each window, we have a triple of numbers (w_i, f_i, s_i) . The set of all those triples for a given user is the representation that we learn. The number of those triples, and the the actual values of each w_i , as per the Temporal Asymmetry Hypothesis, will vary among users and within a given user for different i .

Step 2: Outlier Detection. The personalized representation we derive for each user ((w_i, f_i) in the simplest case) is essentially a set of data points corresponding to every time window. Given a definition of “expected” behavior, we thus can use outlier detection algorithms such as Local Outlier Factor (LOF) [15] that marks points as outliers in a non-parametric manner.

In order to improve LOF’s ability to identify meaningful outliers, instead of processing each user’s set of points/windows individually, we process all users jointly. The rationale behind this choice is that LOF needs a large number of data points that define the “typical” trend; since we define an outlier the same way for each user, and the representation that we learn maps every user into the same (incoming interactions, self-posts) space, using the points from all users can improve outlier detection *via* LOF (in fact, we observed that to be true in early experimentation).

After identifying outlier windows, iNET outputs a set of users that are associated with those windows (there may be cases where more than one outlier windows correspond to a particular user) and highlights the periods for which there was an deviation from the expected behavior.

Case Study: At this step, we argue that the personalized representation (w_i, f_i, s_i) we learn for each user is sufficient for detecting windows of time where the interaction patterns were outlying, according to Definition 2. The reason why this is true is because our representation disassociates the interaction from a physical temporal scale and attaches two feature values (f_i , i.e., the number of incoming interactions, and s_i , the number of self-posts) to each window. These two feature values are thus sufficient to help us detect windows of activity that deviates from the expected pattern.

Furthermore, since our definition of “expected” and “outlying” is based on the ratio of the two features we measure (since Definition 1 defines typical behavior the one for which f_i and s_i are essentially correlated across time), we may treat each window w_i for every user as a data point in the two-dimensional space defined by f_i and s_i and we detect windows of outlying behavior by detecting outliers in this two-dimensional space (i.e., data points / windows for which f_i and s_i are not correlated).

It is important to note that data points that correspond to different windows w_i are observed over widely different time-scales (e.g., one day vs. two months). However, since we are interested in whether f_i and s_i are correlated or not, we are effectively interested in their *ratio* which is a pure number (i.e., it has no units) making joint analysis possible.

Step 3: Outlier Visualization. We present a visualization framework as part of the iNET methodology to understand and analyze outlier intervals that were identified *via* LOF. The specifics of the visualization depend on our case study.

Case Study: For each outlier interval, we provide the word cloud of user text in self posts; by inspection of the word cloud we can identify emerging themes. We leave for future work the incorporation of topic modeling algorithms [16], [17] to the visualization component of iNET. Extracting topical words for the word cloud provides the analyst with a quick overview of

user’s activity while protecting the privacy of the user. Coming back to our Facebook example, we plot the weighted in-degree value (i.e., the network measure of the stable graph in the interval) and the count of self posts on the user’s timeline to create a context around the outlier interval. An example of the visualization framework is shown in Figure 1.

Algorithm 1: iNET Methodology for determining changes in nature of interactions in a user’s timeline

Input: Stream of time labeled edges of user interactions
Number of outliers N

Output: Visualization of outlier intervals activity for top N outliers

// Step 1

- 1: Solve Problem 2 for each user and calculate (w_i, f_i)
- 2: Calculate s_i for the user in each w_i window. For instance, in our case s_i is the number of self posts

// Step 2

- 3: Create Scatter plot where each point corresponds to a window w_i and the two dimensions are s_i and f_i . This includes points for all users
- 4: Determine top N outliers using LOF [15]
- 5: Consolidate the outlier intervals to the respective users

// Step 3

- 6: Visualize outlier intervals as described in Step 3
-

B. Discussion

Handling Streaming Data: iNET operates on streaming data of user interactions and does not require the entire timeline of a user in order to produce a set of anomalous results.

Furthermore, iNET is an *anytime algorithm* in the sense that it can produce an output at any given point in time which will reflect the current state of the personalized interaction networks, as they have been observed in the stream. This is important because this implies that iNET can be deployed and track changes in users’ nature of interaction in near real-time for the portion of the stream it has observed so far.

Sensitivity to User Privacy: In order for iNET to work, we do not require access to the content (e.g., text or photos) that a user or their friends share. Such content is very frequently personal and private in nature and users may feel uncomfortable with sharing such private information for the sake of our analysis. Furthermore the representation that iNET learns requires only aggregate statistics of a user’s interaction network, without keeping track of the particular individuals with whom the user is interacting. If, for instance, data collection is done via a custom-built social network application, the entire functionality Step 1 of iNET can be pushed to the application and the only information accessible to the analyst would be the aggregate statistics of the personalized representation.

As part of evaluating iNET, in Section V we use the text posted by the user during a given outlier interval to analyze the context and understand the reasoning as to what has given rise to the outlier indicating change in the nature of user

interactions. This, however, is done with the user’s consent and is strictly for evaluation purposes. We believe that iNET is minimally invasive to a user’s privacy and is a step towards exploring the trade off of how invasive such a method needs to be in order to effectively identify anomalous events.

Generality: iNET is a general methodology that can be applied to a variety of time-evolving user interaction networks. In our particular realization of iNET, we use Facebook as our case study and our running example, and we tailor our notions of “expected” and “outlying” with respect to that example (see Definitions 1 and 2). However, the methodology is directly applicable to other scenarios, as long as the features that are used in lieu of f_i and s_i are aggregated in the same unit, or can be easily transformed to the same unit. For instance, in our case f_i and s_i are the number of incoming and self posts respectively over the course defined by window w_i and their units are “posts per duration of window w_i , and their ratio is a pure number, which enables us to compare the behavior across different windows. We may also add more features that we measure per window and in that case, the outlier detection step would be carried out in the space defined by those features. For example, if we observe interactions between the peers of the center of our star graph (i.e., the user whose personalized representation we are learning), we can use additional graph metrics such as the number of triangles in our analysis.

V. EXPERIMENTS

In this section, we present the experiments using iNET on a Facebook user dataset. We discuss results of outliers obtained using variable-length intervals and compare and contrast with those outliers obtained from fixed-length intervals.

A. D2y Dataset

The dataset *D2y* is an 18.5-month longitudinal Facebook dataset collected *via* an application voluntarily installed by Facebook users. It consists of 831 Facebook users and all the activity on their Facebook walls. The duration of the dataset is approximately 18.5 months, 13 Apr 2012 through 30 Oct 2013. A user’s Facebook wall/timeline can be described as their space where users express their ideas and thoughts, and interact with other users through likes, comments and status updates. We analyzed 381,690 wallposts out of which 295,265 were authored by the users themselves, which we call *self posts* and the 86,425 were posted on the users’ walls by their friends as *friend posts*. We extracted 415,589 interactions between users and their friends occurring in the form of likes, comments and posts, from the walls of users. Each interaction is represented by a temporal edge (from-id, to-id, timestamp), where from-id is the source of the interaction, to-id is the user whom the interaction is directed towards and the timestamp pertains to when the interaction occurred.

B. Using iNET on D2y dataset

We used iNET to analyze users in the *D2y* dataset described above. We use weighted in-degree network statistic to compute

variable-length intervals using Step 1 of our iNET methodology. At each aggregation step, it measures the weighted in-degree of the aggregated edges until it meets the stopping criterion and repeats the interval convergence process until there are no more temporal edges left in the user’s interaction stream. Once the variable-length intervals are computed for each user, we obtain the number of self posts that the user has made in these interval periods. We combine the weighted in-degree network statistic and the self posts of all users in all intervals for efficiency and compute the top N outliers (we considered the top 30) using the LOF method and analyze those outlier windows using the visualization framework. Some of the example outliers of our analysis can be observed in Figure 4.

The visualization framework presents plots of the weighted in-degree and the count of self posts in each aggregation interval on the timeline of the user. This helps understand how those interaction values varied from interval to interval in the user’s timeline.

Exclusively for validation. Although we expressly avoid analyzing the text during the identification of the intervals (to ensure that our method is minimally invasive to user privacy), we look at the text for the purposes of interpreting the results. We present a word cloud from the text in user’s self posts in the outlier interval to infer what external event has given rise to an outlier interval. Figure 4 presents the visualization of four different users with outlier intervals identified in an unsupervised fashion by iNET.

The outlier intervals are classified into nine categories (shown in Table I) based on the frequency of words that appeared in the corresponding user posts. We argue that this approach not only presents a quick overview of the user’s activity, but also protects the privacy of user. An interval can belong to more than one category. The word cloud in three outlier intervals of user #13 in Figure 4(a) depicts that the user shares a lot of posts related to Politics, where the words ‘Obama’, ‘new’, and ‘President’ are repeated in every outlier interval. In addition, we observe that the number of self posts in the outlier interval is slightly higher than in the previous interval, which could suggest that the occurrence of an external event has triggered the user towards sharing more posts than usual. This trend can also be observed with weighted in-degree plot presented in the top portion of the visualization, indicating that the user’s friends engage more than usual with the user during the outlier interval *via* likes, comments, and posting on user’s wall. Further inquiry into user’s sharing behavior has shown us that the user shares a lot of link-type posts from news websites. From the words in the user #13’s outlier interval word cloud, we categorize the user’s behavior as interested in Politics. Figure 4(b) only has four variable-length intervals out of which the outlier interval has more number of interactions than usual, but lower number of self posts. The top words in the outlier interval word cloud are ‘Apple’, ‘new’, ‘iPhone’, ‘company’ and ‘cloud’. We categorize the user’s interest during the outlier period as Technology and News. Figure 4(c) presents the visualization for User #22.

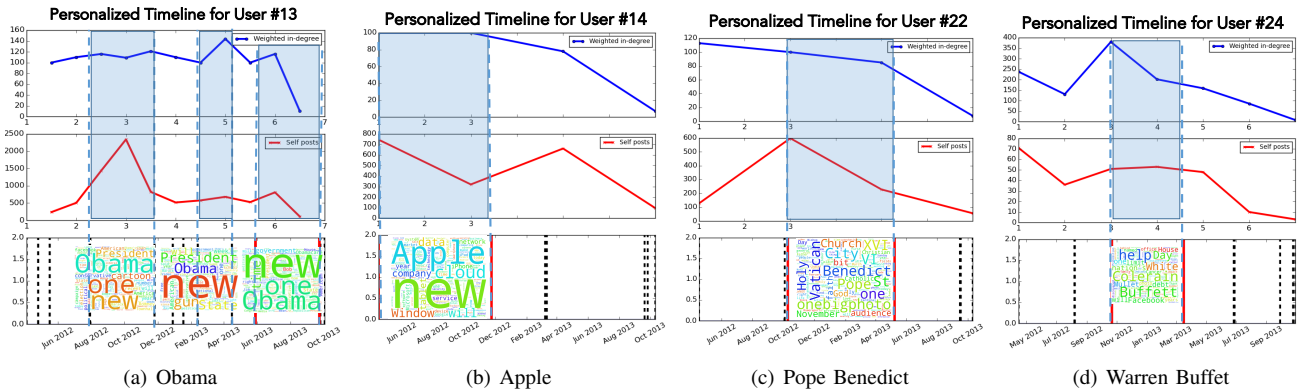


Fig. 4. Four indicative results of the Visualization component of iNET framework. By exception, we also include the word clouds for each interval of interest in order to enable the interpretation and validation of the results.

Category	iNET	iNET-fixed			
		24-hour	1-week	1-month	3-month
Politics	13	15	17	15	14
Personal	11	0	9	10	4
News	8	0	4	3	7
Religion	4	0	0	0	0
Movies	2	0	0	0	0
Music	4	0	0	3	2
Sports	3	0	0	2	2
Technology	2	0	0	0	0
Misc	2	15	7	3	6

TABLE I

CATEGORIZATION OF VARIABLE-WINDOW AND FIXED-WINDOW OUTLIERS USING CONTENT IN THE OUTLIER INTERVAL. THE HIGHLIGHTED ROWS INDICATE CATEGORIES THAT ARE OF LOW INTEREST SINCE WE WERE EITHER UNABLE TO LABEL, OR THEY CORRESPOND TO “ORGANIC” EVENTS THAT HAVE A LIFE-SPAN OF DAY, SUCH AS BIRTHDAY CELEBRATIONS. WE OBSERVE THAT iNET OFFERS A WIDE VARIETY OF NON-ORGANIC AND RELEVANT EVENTS, COMPARED TO iNET-FIXED WHICH IS RESTRICTED TO A SINGLE-DAY OBSERVATION WINDOW.

The user has both high number of interactions as well as self posts during that period. Words like ‘Pope’, ‘Benedict’, ‘Holy’, ‘Vatican’, and ‘God’ prominently appear in the word cloud during the outlier period and we categorize the user’s interest as Religion. In the Figure 4(d), the number of interactions as well as self posts in the outlier interval is lower than the previous interval and words like ‘Colerain’, ‘Buffett’, ‘Mullet’, ‘White’ and ‘House’ appear prominently in the word cloud of the user. Given the user’s penchant to share information, we label this user #22’s outlier interval as News and Politics. Notice the distinct number of variable-length intervals in each user in Figure 4 which is further validation of our Temporal Asymmetry Hypothesis.

C. What about using fixed-length intervals?

iNET computes variable-length intervals depending on the rate of interactions in the user stream and outputs the visualization of user’s activity in the outlier interval to help the analyst understand the context of the outlier. In order to see how well iNET performs against using fixed-length intervals, we introduce iNET-fixed with four different durations of 24-hour, 1-week, 1-month, and 3-months, as a baseline approach.

iNET-fixed uses the same methodology as iNET, except for the step where it computes variable-length intervals. For example, the duration of the D2y dataset is 566 days, which gives 566 intervals of fixed-length, each with a duration of 24 hours. In like manner, D2y dataset can be partitioned into 82 intervals of 1-week, 19 intervals of 1-month and 6 intervals of 3-month duration each. For a user, for each fixed interval, we aggregate all incoming interactions into a graph and compute the weighted in-degree network statistic for that period. In addition, we get the count of self-posts on the user’s wall for a given fixed interval. For the purposes of efficiency, we compute outlier scores for the pairs of values (weighted in-degree, self-posts) using LOF and output the visualization framework for the top 30 outliers using iNET-fixed. Since we manually label the categories of detected outliers, we chose 30 outliers, which is a large enough number for evaluation purposes, yet manageable for labeling.

Our visualization framework provides a quick way for the analyst to understand the context of the outlier window. We classify the top 30 outlier intervals obtained *via* iNET and the iNET-fixed into categories, which we present in Table I. The categories *Politics* and *Personal* seem highly prevailing in both the methodologies. The *Personal* category consists of personal things occurring in the interval including the names of user’s friends and discussions that are only pertinent to user’s life. As an indicative example, user #10’s outlier interval has words such as ‘tripbday’, ‘memorialbbqcalm’, ‘chowder’, ‘plan’, ‘family’ and ‘drunk’, along with name mentions, which we believe are user’s friends. We do not present the related visualization to maintain user anonymity. Figure 4(a) presents outlier windows belonging to the category *Politics*, where the word cloud contains all the words related to American politics. In addition, we observe categories *News*, *Religion*, *Movies*, *Music*, *Sports*, and *Technology*. The category *Misc* are those outlier intervals that did not exhibit a thematic coherence and user birthdays that occur without any interference from the user.

a. iNET discovers more diverse and relevant categories compared to iNET-fixed. We computed the Shannon’s entropy values for the two methodologies to measure the di-

versity in outlier intervals. Intuitively, if a method identifies only one category, its entropy is zero. Higher entropy reflects higher diversity in the discovered outliers and we observe this with iNET. iNET-fixed has entropy of 0.76, 1.88, 2.1, and 2.4 respectively for durations, 24-hour, 1-week, 1-month, and 3-months, where iNET has the highest entropy with 2.75 bits. Note that we excluded the Misc category from the precision calculation, since Misc does not represent a single thematically-coherent category.

b. Variable size windows identify meaningful user “phases” more effectively. Going a step deeper, we examined the categories of the windows adjacent (previous and next) to the discovered outliers for both iNET and iNET-fixed. In the majority of cases in iNET-fixed, there was no change in category from the adjacent intervals. For example, if the outlier interval with a 1-week duration is of category Politics, the adjacent intervals usually belonged the same category. However, with iNET, we observed that the category of intervals adjacent varied from that of the category of the discovered outlier. iNET captured 100% of relevant categories with a 95% precision, compared to the baselines.

c. Exogenous influence on user behavior: Based on the word cloud of the outlier interval of user #8, we categorized it as Politics, where the top words are ‘India’, ‘Gujarat’, ‘Narendra’, ‘Modi’, ‘leadership’ and ‘visionary’. The duration of the interval is 126 days with period being Jun-Sep 2013. The time period coincides exactly with the elections in India, where a politician declared running for prime minister’s office in September 2013. The categories of the subsequent windows are Personal and Sports, which differ from that of the outlier interval category, further confirming our belief that variable-length intervals provide for better observation of user behavior.

iNET provides better understanding towards identifying change in behavior since the overarching goal is to identify those periods of time for which there was a significant deviation of a user’s nature of interaction without any supervision, letting the data answer the question for us. By placing the constraint on the iNET-fixed that the intervals be of a certain size, we *impose* a certain kind of supervision on the algorithm towards finding the outlier periods in the user’s timeline. Of course, selecting the network measure f_i (e.g., the weighted in-degree) introduced a bias to our analysis, however, we believe that such a bias is much more flexible than committing to a single, fixed-size interval.

D. Discussion

Flexibility of observation intervals: While observing underlying intervals in users, we observed that there were some internal and external factors that influence user behavior. For example, users interacted more during the 2012 US elections, sharing opinions and posting news. In addition, the adjacent intervals belonged to different category than that of the outlier, which was also the case in majority of the outlier intervals discovered *via* iNET. While change in behavior based on external events is not surprising, fixed-length intervals do not provide any information about change-points without some

kind of manual supervision. Based on the examples provided earlier, we observed that iNET detects change-points as well as the beginning and end of these change points, thus providing us the with appropriate observation intervals determined by user interaction patterns.

Similarity search: One of the applications for iNET is finding users similar in their interaction signature irrespective of the timescale. In order to simplify the representation that iNET learns, we may take the ratio of the interaction features in each interval to present a concise, unit-free interaction signature. Each personalized representation now is essentially a time series where the time axis is defined by w_i . Thus, given a user of interest who was indicated to us by Step 2, we can search for other users in our collection of personalized representations that share the same interaction signature over time. It is important that we not restrict our search to cases where the interaction signatures perfectly align with respect to the (arbitrary) time axes w_i . We should allow for cases where the same pattern is observed in “elongated” or “shorter” forms, but the overall trend over those arbitrary windows of time is the same. For that purpose we can use the so-called Dynamic Time Warping (DTW) distance [18] for time-series. Once we find outliers using iNET, we can employ DTW to find users with similar interaction patterns as outliers. We reserve investigation of this application for future work.

VI. RELATED WORK

Our work lies in the intersection of various topics that researchers have extensively studied in the past. We discuss the state of the art for each of those areas in the next few lines. However, we must point out that to the best of our knowledge, our viewpoint of conducting such an analysis on a personalized, per-user basis and introducing the Temporal Asymmetry Hypothesis that is instrumental in the derivation of iNET has not been studied before.

A. Time-evolving networks and granularity

A key characteristic of dynamic interaction networks is their continual change. Researchers have developed systematic approaches on partitioning a data stream. Sulo et al. [19] identify appropriate aggregation intervals by balancing the trade-off between smoothness and noise in the network. [12] proposed DAPPER, which uses a window size and a shift parameter to calculate a frequency change vector and partition data based on changes in network structure. GRAPHSCOPE [20] is an MDL-based, parameter-free algorithm that merges “similar” snapshots into a segment and compresses them together; on the other hand, “dissimilar” consecutive snapshots lead to the creation of a new segment, and declaration of a change-point. [21] proposed a supervised time scale detection framework that leverages ground truth from training data for “good windowing” based on the task at hand, predicting links, attributes, or change-points. These works identify fixed-length intervals and require the entire data stream at once. Several previous works address problems like finding interesting patterns [22], mining periodic behaviors [23] as well as subgraphs [24] from

longitudinal networks. For our work, we use ADAGE [14], which is an online approach for partitioning streaming graph data into variable-length intervals. While these works focused on evolution of dynamic global graphs, we focus on individual dynamic interaction graphs to analyze user behaviors in time.

B. Anomaly detection

Anomaly or outlier detection is a major direction in the data mining community and there exist several surveys for a variety of static and time-evolving data [25], [26], including graphs [27].

Clouds of points. For outliers in clouds of points the representative algorithm is the Local Outlier Factor (LOF) [28], which marks points as outliers if their local density is *different* from the density of their neighbors. Improved techniques include LOCI [29] that detects outliers and micro-clusters without user-defined parameters, and COF [30] that distinguishes the cases of isolation and low density. OddBall [1] detects anomalies in *static* graphs by focusing on egonet-related properties, depicting pairs of them in plots, and leveraging LOF-based outlieriness scores to detect original nodes that are anomalous in the projected space. OPAVion [31] combines feature aggregation, outlier detection (via OddBall) and visualization (via Apollo [32]) into a system that automatically mines and interactively visualizes static anomalies in large graphs.

Change point detection. Work in this space attempts to identify points in time when the observed network (or some substructures) change significantly. Clustering is a common technique used to that end: [33] aims at change detection in streaming graphs using projected clustering; Com2 [34] uses graph-search and PARAFAC tensor decomposition followed by MDL to find dense temporal cliques and bipartite cores. [35] uses incremental cross-association for change detection in dense blocks over time.

VII. CONCLUSIONS

We introduce iNET, a novel analytic framework, in the intersection of dynamic and time-evolving graph mining, anomaly detection, and personalization on the web, and is a systematic approach for identifying changes in users' mode of interaction. Indicatively, in our case study, iNET produces results that span the entire spectrum of categories while the best performing baseline covers only 66%. We view this paper as first stepping stone towards exploring the trade-off between the minimum invasive information needed to detect behavioral change, and our ability to relate online behavioral changes to real-life changes of the user.

REFERENCES

- [1] L. Akoglu, M. McGlohon, and C. Faloutsos, "OddBall: Spotting anomalies in weighted graphs," *Advances in Knowledge Discovery and Data Mining*, pp. 410–421, 2010.
- [2] P. Devineni, D. Koutra, M. Faloutsos, and C. Faloutsos, "If walls could talk: Patterns and anomalies in facebook wallposts," in *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015*, ser. ASONAM '15. New York, NY, USA: ACM, 2015, pp. 367–374. [Online]. Available: <http://doi.acm.org/10.1145/2808797.2808880>
- [3] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee, "Measurement and analysis of online social networks," in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*. ACM, 2007, pp. 29–42.
- [4] J. Ugander, B. Karrer, L. Backstrom, and C. Marlow, "The anatomy of the facebook social graph," *arXiv preprint arXiv:1111.4503*, 2011.
- [5] M. De Choudhury, S. Counts, E. J. Horvitz, and A. Hoff, "Characterizing and predicting postpartum depression from shared facebook data," in *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, ser. CSCW '14. New York, NY, USA: ACM, 2014, pp. 626–638.
- [6] E. L. Murnane and S. Counts, "Unraveling abstinence and relapse: Smoking cessation reflected in social media," in *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*, ser. CHI '14. New York, NY, USA: ACM, 2014, pp. 1345–1354. [Online]. Available: <http://doi.acm.org/10.1145/2556288.2557145>
- [7] CALO, "Enron email network," 2009. [Online]. Available: <http://www.cs.cmu.edu/enron/>
- [8] N. Koblenz, <http://konect.uni-koblenz.de/networks/>, 2016, accessed: Dec-2016.
- [9] B. Viswanath, A. Mislove, M. Cha, and K. P. Gummadi, "On the evolution of user interaction in facebook," in *Proceedings of the 2nd ACM Workshop on Online Social Networks*, ser. WOSN '09. New York, NY, USA: ACM, 2009, pp. 37–42.
- [10] D. Koutra, E. E. Papalexakis, and C. Faloutsos, "TensorSplat: Spotting Latent Anomalies in Time," in *Informatics (PCI), 2012 16th Panhellenic Conference on*. IEEE, 2012, pp. 144–149.
- [11] L. Akoglu and C. Faloutsos, "Event detection in time series of mobile communication graphs," in *Army Science Conference*, 2010, pp. 77–79.
- [12] R. S. Caceres, "Temporal scale of dynamic networks," Ph.D. dissertation, University of Illinois at Chicago, 2013.
- [13] J. Kiernan and E. Terzi, "Constructing comprehensive summaries of large event sequences," *ACM Trans. Knowl. Discov. Data*, vol. 3, no. 4, pp. 21:1–21:31, Dec. 2009.
- [14] S. Soundarajan, A. Tamersoy, E. B. Khalil, T. Eliassi-Rad, D. H. Chau, B. Gallagher, and K. Roundy, "Generating graph snapshots from streaming edge data," in *Proceedings of the 25th International Conference Companion on World Wide Web*, ser. WWW '16 Companion. Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee, 2016, pp. 109–110. [Online]. Available: <http://dx.doi.org/10.1145/2872518.2889398>
- [15] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "Lof: identifying density-based local outliers," in *ACM sigmod record*, vol. 29. ACM, 2000, pp. 93–104.
- [16] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman, "Indexing by Latent Semantic Analysis," *Journal of the American Society for Information Science*, vol. 41, no. 6, pp. 391–407, Sep. 1990.
- [17] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *Journal of machine Learning research*, vol. 3, no. Jan, pp. 993–1022, 2003.
- [18] E. Keogh, "Exact indexing of dynamic time warping," in *Proceedings of the 28th International Conference on Very Large Data Bases*, ser. VLDB '02. VLDB Endowment, 2002, pp. 406–417. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1287369.1287405>
- [19] R. Sulo, T. Berger-Wolf, and R. Grossman, "Meaningful selection of temporal resolution for dynamic networks," in *Proceedings of the Eighth Workshop on Mining and Learning with Graphs*, ser. MLG '10. New York, NY, USA: ACM, 2010, pp. 127–136. [Online]. Available: <http://doi.acm.org/10.1145/1830252.1830269>
- [20] J. Sun, C. Faloutsos, S. Papadimitriou, and P. S. Yu, "GraphScope: parameter-free mining of large time-evolving graphs," in *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, ser. KDD. ACM, 2007, pp. 687–696.
- [21] B. Fish and R. S. Caceres, "A supervised approach to time scale detection in dynamic networks," *arXiv.org*, Feb. 2017. [Online]. Available: <https://arxiv.org/abs/1702.07752>
- [22] J. Leskovec, J. Kleinberg, and C. Faloutsos, "Graph evolution: Densefication and shrinking diameters," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 1, no. 1, p. 2, 2007.
- [23] M. Lahiri and T. Y. Berger-Wolf, "Mining periodic behavior in dynamic social networks," in *2008 Eighth IEEE International Conference on Data Mining*. IEEE, 2008, pp. 373–382.

- [24] P. Bogdanov, M. Mongiovì, and A. K. Singh, "Mining heavy subgraphs in time-evolving networks," in *2011 IEEE 11th International Conference on Data Mining*. IEEE, 2011, pp. 81–90.
- [25] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Comput. Surv.*, vol. 41, no. 3, 2009.
- [26] A. B. Varun Chandola and V. Kumar, "Anomaly Detection for Discrete Sequences: A Survey," *IEEE TKDE*, vol. 24, pp. 823–839, 2012.
- [27] L. Akoglu, H. Tong, and D. Koutra, "Graph-based Anomaly Detection and Description: A Survey," *Data Mining and Knowledge Discovery (DAMI)*, April 2014.
- [28] M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "LOF: Identifying Density-Based Local Outliers," in *SIGMOD*, 2000.
- [29] S. Papadimitriou, H. Kitagawa, P. B. Gibbons, and C. Faloutsos, "LOCI: Fast Outlier Detection Using the Local Correlation Integral," in *ICDE*, 2003, pp. 315–326.
- [30] J. Tang, Z. Chen, A. W. Fu, and D. W. Cheung, "Enhancing Effectiveness of Outlier Detections for Low Density Patterns," in *PAKDD*, 2002, pp. 535–548.
- [31] L. Akoglu*, D. H. Chau*, U. Kang*, D. Koutra*, and C. Faloutsos, "OPAvion: Mining and Visualization in Large Graphs." ACM, 2012, pp. 717–720.
- [32] D. H. Chau, A. Kittur, J. I. Hong, and C. Faloutsos, "Apolo: Making Sense of Large Network Data by Combining Rich User Interaction and Machine Learning," in *KDD*, 2011.
- [33] C. C. Aggarwal and S. Y. Philip, "Online Analysis of Community Evolution in Data Streams," in *SDM*. SIAM, 2005.
- [34] M. Araujo, S. Papadimitriou, S. Günemann, C. Faloutsos, P. Basu, A. Swami, E. E. Papalexakis, and D. Koutra, "Com2: Fast Automatic Discovery of Temporal ("Comet") Communities." Springer, 2014, pp. 271–283.
- [35] J. Ferlez, C. Faloutsos, J. Leskovec, D. Mladenic, and M. Grobeldnik, "Monitoring Network Evolution using MDL," *ICDE*, pp. 1328–1330, 2008.