

Discovering topically- and temporally-coherent events in interaction networks

Han Xiao¹, Polina Rozenshtein², and Aristides Gionis²

¹Helsinki Institute for Information Technology and
Department of Computer Science
University of Helsinki, Finland
`hxiao@cs.helsinki.fi`

²Helsinki Institute for Information Technology and
Department of Computer Science
Aalto University, Finland
`firstname.lastname@aalto.fi`

Abstract. With the increasing use of online communication platforms, such as email, Twitter, and messaging applications, we are faced with a growing amount of data that combine *content* (what is said), *time* (when), and *user* (by whom) information. Discovering meaningful patterns and understand what is happening in this data is an important challenge. We consider the problem of mining online communication data and finding top- k *temporal events*. A temporal event is a coherent topic that is discussed frequently in a relatively short time span, while its information flow respects the underlying network.

Our method consists of two steps. We first introduce the notion of *interaction meta-graph*, which connects associated interactions. Using this notion, we define a *temporal event* to be a subset of interactions that (i) are topically and temporally close and (ii) correspond to a tree that captures the information flow. Finding the best temporal event leads to a budget version of the prize-collecting Steiner-tree (PCST) problem, which we solve using three different methods: a greedy approach, a dynamic-programming algorithm, and an adaptation to an existing approximation algorithm. Finding the top- k events maps to a maximum set-cover problem, and thus, solved by greedy algorithm. We compare and analyze our algorithms in both synthetic and real datasets, such as Twitter and email communication. The results show that our methods are able to detect meaningful temporal events.

Keywords: social-network analysis, temporal networks, event detection

1 Introduction

Event detection is a fundamental data-mining problem in many different domains, such as, time series and data streams [10], point clouds and vector spaces [4], and networks [3]. In this paper we focus on the problem of detecting

events in networks, in particular, networks that contain both *content* and *time* information. An *interaction* (u, v, α, t) occurs whenever a piece of information α is exchanged between two network entities u and v at time t . Examples of interaction networks include data communication networks, such as email, Twitter, or online messaging systems.

Our goal is to summarize the network activity by finding the top- k events. We consider an interaction network $H = (N, \mathcal{I})$, where interactions \mathcal{I} take place among a set of network entities N . The interactions in \mathcal{I} are directed, annotated with content information, and time-stamped. We define an *event* in the interaction graph H to be a subset of interactions, $\mathcal{I}' \subseteq \mathcal{I}$ that are (i) temporally close, (ii) topically similar, and (iii) correspond to a tree that captures the information flow in the network. The intuition behind representing events as trees is similar to the work by Yang [19].

We convert the interaction network $H = (N, \mathcal{I})$ into a weighted *interaction meta-graph* $G = (\mathcal{I}, E)$, that is, a graph whose vertices are the interactions \mathcal{I} . Two interactions $i, j \in \mathcal{I}$ are connected in G if it is possible to explain the information flow between i and j . In particular, we consider three types of flow: *broadcast*, *relay* and *reply*. The edge weights of the interaction meta-graph G measure the topic dissimilarity between connected interactions. Our transformation from the interaction network to the interaction meta-graph has the interesting property that an *event* in the interaction graph H corresponds to a tree T in the interaction meta-graph G . The root of the tree T is interpreted as the source of the event. Downstream interactions (interactions that are reachable from the root) are due to information propagation.

Motivated by the previous discussion, we formalize the task of interaction-network summarization as the problem of finding top- k trees in the transformed interaction meta-graph $G = (\mathcal{I}, E)$. We decompose this task into two sub-problems. First, we find a set of independent candidate events that are temporally and topically coherent. Since our goal is to summarize the interaction network we aim to find large events. We show that this problem is the budget version of prize-collecting Steiner-tree problem in directed acyclic graphs. We provide three algorithms, among which a greedy approach performs the best.

The second sub-problem is to select k events that maximize the overall node coverage. This task maps to the maximum set-cover problem, and it can be approximated using a standard greedy algorithm. To speed up further our algorithm, we also propose a search strategy that avoids evaluating candidate events at all possible tree roots, but heuristically selects the most promising ones.

Example 1. Consider the email communication network of a company, such as the one shown in Fig. 1. The interaction network is shown in Fig. 1(a) and the corresponding interaction meta-graph in Fig. 1(b). The edges between interactions $(2, 4)$, $(1, 2)$, $(2, 3)$ in Fig. 1(b) are examples of edge types broadcast, relay, and reply, respectively. In this toy example there are two main events. (i) *progress*: The CEO asks a project manager (PM) about progress on a project, and the PM forwards the request to team members 1 (TM1) and 2 (TM2). Later, TM1 reports back to PM, who in turn reports back to CEO. The information flow of

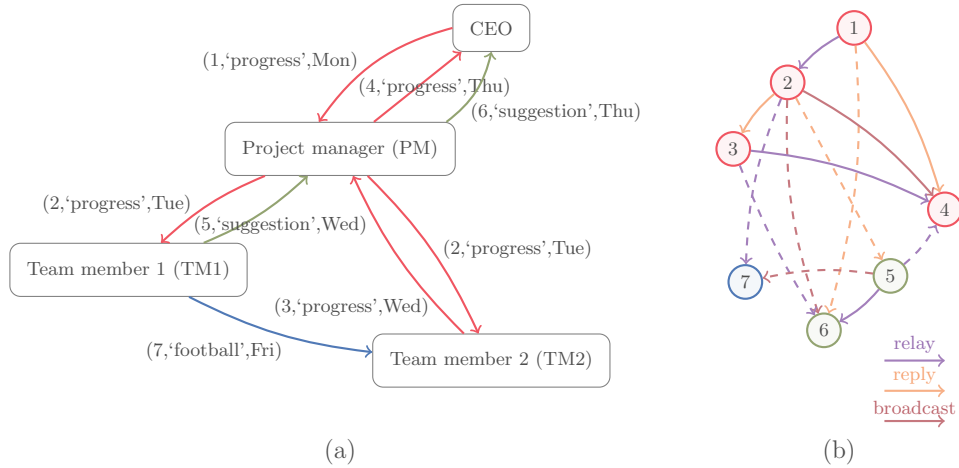


Fig. 1: A toy example showing the email communication network within a company. (a) The interaction network. Each edge corresponds to one interaction/email, labeled as (**interaction id**, **message topic**, **timestamp**). (b) The corresponding interaction meta-graph. Topics in both graphs are depicted using different colors. Edges in the interaction meta-graph are depicted by a different color according to their type (relay, reply, or broadcast). Edges are solid if they have small weight (topic dissimilarity). Otherwise, they are dashed.

this event follows the interactions $1 \rightarrow 2 \rightarrow 3 \rightarrow 4$. (ii) *suggestion*: Motivated by the first event, TM2 comes up with some suggestion, which she sends to PM. The PM finds the suggestion useful and forwards it to CEO. The information flow of this second event is $5 \rightarrow 6$. A third event, *football*, is smaller in size, and it is not included in the top-2 events. Note that due to time ordering of the interactions, the interaction meta-graph G is a directed acyclic graph.

The problem considered in this paper has many applications in different domains. In our experimental evaluation, we focus on analyzing textual data in social media. We experiment with one email dataset (Enron) and three Twitter datasets. We provide a comparison of the different approaches, as well as many examples in which our methods discover meaningful events.

The contributions of this paper are summarized as follows¹:

- We propose a novel formulation for the problem of discovering events that are temporally and topically coherent in interaction networks, such as, online communication networks.
- We present a transformation of the interaction network to an interaction meta-graph, which captures temporal and topical association of interactions

¹ All scripts are available at <https://github.com/xiaohan2012/1st>

as well as the information flow in the network. This transformation helps to provide a cleaner abstraction to the event-detection problem.

- For the problem of finding high-volume events while satisfying constraints of temporal and topical coherence we present and we evaluate three different algorithms: a greedy approach, a dynamic-programming algorithm, and an adaptation to an existing approximation algorithm.
- We address the problem of finding the top- k events that summarize the network activity. The classic greedy algorithm is the standard way to approach this problem, but here, to speed-up the computations, we also propose and evaluate a search strategy that avoids construction of candidate events at all possible tree roots, but adaptively selects the most promising ones.
- We compare and analyze our algorithms on both synthetic and real datasets, such as Twitter and email communication. We show that our methods are able to detect meaningful temporal events.

2 Related work

Phrase-based event detection. The problem of detecting events in social media has attracted significant attention. Leskovec et al. [14] and Yang et al. [20] treat events as short, distinctive phrases that propagate relatively intact through in a network. Their work offers a graph formulation for clustering variants of phrases based on string edit distance. Although their objective is similar to ours, there are significant differences. First, our methods focus on interaction networks, aiming to capture information flows in communication networks, rather than action networks. Second, we explicitly impose topic-coherence constraints, where the edit distance is insufficient for this goal. Third, instead of representing events by phrases, we derive higher-level representation using topic terms.

Text summarization. Text summarization techniques attempt to select a subset of sentences [6] or tweets [11] to summarize textual content. Similarly, we select a subset of interactions under a topic-coherence constraint. However, we also impose temporal coherence constraint, whereas they take a static view.

Statistical methods. Statistical and machine learning approaches for event detection are gaining increasing attention in recent years. Mathioudakis et al. [15] develop an interactive system for identifying trends (events). The system first identifies “bursty” keywords, then clusters them based on co-occurrence and later performs trend analysis using dimension-reduction methods. Becker et al. [1] focus on online event identification. Their approach relies on online clustering techniques in order to discover topically-related tweets as an event and feature-based modeling in order to distinguish events from non-events. The difference of this approach with our work is that we offer a graph-theoretic formulation.

Graph-based methods. Other event-detection methods are based on constructing a word graph [5, 16, 18]. Weng et al. [18] combines wavelet analysis and graph-partitioning techniques to cluster the words into events. Meladianos et al. [16] construct a word graph to represent a sequence of tweets, however,

they focus on identifying key sub-events inside the sequence. Cataldi et al. [5] detect events by locating strongly connected components. Compared to those approaches, in this paper we explicitly model interactions, and take into account temporal constraints and topical-coherence constraints.

3 Model

An *interaction network* $H = (N, \mathcal{I})$ consists of a set of n nodes N and a set of m time-stamped interactions \mathcal{I} between pairs of nodes. \mathcal{I} is represented as:

$$\mathcal{I} = \{(u_i, v_i, \alpha_i, t_i)\}, \text{ with } i = 1, \dots, m, \text{ such that } u_i, v_i \in N, t_i \in \mathbb{R}, \alpha_i \in \mathbb{R}^L,$$

indicating that nodes u_i and v_i interacted at time t_i . Each interaction is annotated with textual content represented by α_i . The representation is independent to our main methodology. We can use various text modeling techniques such as bag-of-words representation or latent Dirichlet allocation (LDA) [2].

For generality we consider that interactions are directed. More than one interaction may take place between a pair of nodes, with different timestamps. Conversely, more than one interaction may take place at the same time, between different nodes. Online communication networks, such as email networks, are examples of interaction networks.

Given an interaction network H we construct a directed weighted *interaction meta-graph* $G = (\mathcal{I}, E, c)$. The vertices \mathcal{I} in G correspond to the interactions \mathcal{I} in H . There is an edge from vertex $i = (u_i, v_i, \alpha_i, t_i) \in \mathcal{I}$ to a vertex $j = (u_j, v_j, \alpha_j, t_j) \in \mathcal{I}$ if the following holds:

1. Interaction i takes place before interaction j (time comprehension): $t_i \leq t_j$.
2. Information comprehension takes place in one of the following ways:
 - (a) interactions i and j share the same start node in N : $u_i = u_j$ (broadcast);
 - (b) the end node of interaction i is the start node of interaction j and the end node of j is not the start node of i : $v_i = u_j$ and $v_j \neq u_i$ (relay);
 - (c) the end node of an interaction i is the start node of an interaction j and the end node of j is the start node of i : $v_i = u_j$ and $v_j = u_i$ (reply).

Note, that due time comprehension the G is a directed acyclic graph (DAG).

For the edges of the interaction meta-graph G we use weights to measure the topical (dis)similarity between interactions. Thus, given two interactions $(u_i, v_i, \alpha_i, t_i)$ and $(u_j, v_j, \alpha_j, t_j)$ connected by an edge in G , our edge-weighting function $c: E \rightarrow \mathbb{R}$ is a distance function between topic vectors α_i and α_j .

Finally, given a meta-graph $G = (\mathcal{I}, E, c)$ and a time interval $[s, f]$ we define the *time-induced meta-graph* $G([s, f]) = (\mathcal{I}([s, f]), E, c)$, where $\mathcal{I}([s, f])$ are the interactions that occur in $[s, f]$: $\mathcal{I}([s, f]) = \{(u, v, \alpha, t) \in \mathcal{I} \mid s \leq t \leq f\}$.

4 Problem formulation

We aim at summarizing the top- k events in an interaction network. We define an *event* to be a rooted subtree T of the interaction meta-graph G . An event naturally has a source vertex (or interaction) and is spread in the network. We are interested in events of high volume, which translates into a large number of iterations included into the tree T . We are also interested in events with temporally close and topically coherent interactions.

These aspects can be incorporated into the optimization cost function in different ways. Our primary objective is to obtain k events that have high enough coverage to represent the whole network, and thus, we aim to maximize the number of interactions that are included in the event. To incorporate temporal and topical coherence we set constraints on the time interval spanned by the event tree (temporal coherence), and on total weight of its edges (topical coherence).

To simplify the problem of finding the best k events, we decompose the main task into two subproblems: (1) finding a set of independent candidate events that satisfy the constraints and maximize volume of interactions, and (2) selecting the top- k events to maximize total coverage. The first problem is defined as follows.

Problem 1. Time-constrained maximum tree (TMAXTREE): Given an interaction meta-graph $G = (\mathcal{I}, E, c)$, a root vertex $r \in \mathcal{I}$, time budget I , and dissimilarity budget B , find a directed subtree $T = (V_e, E_e) \subseteq G$, rooted at r , which satisfies the constraints

$$\sum_{e \in E_e} c(e) \leq B \quad \text{and} \quad (\max_{i \in V_e} t_i - \min_{j \in V_e} t_j) \leq I,$$

while maximizing the number of vertices $|V_e|$.

Note that the time constraint can be omitted, if we restrict the input graph to be induced by the time interval $[t_r, t_r + I]$, where t_r is the root timestamp. By omitting the time constraint, our problem can be written as follows.

Problem 2. Maximum tree (MAXTREE): Given a weighted directed acyclic graph $G([s, f]) = (\mathcal{I}([s, f]), E, c)$, a root vertex r , and cost budget B , find a subtree $T = (V_e, E_e) \subseteq G([s, f])$, rooted at r , that satisfies $\sum_{e \in E_e} c(e) \leq B$ while maximizing the number of vertices $|V_e|$.

We observe that MAXTREE is directly related to *budget* version of the prize-collecting Steiner-tree problem (PCST) [12]. However, we are dealing with a special case of the budget PCST, as vertex prize is uniform and our input graph is a DAG. Despite so, this special case is still **NP**-hard.

Proposition 1. MAXTREE is **NP**-hard.²

² The full version of the paper can be found at <http://arxiv.org/abs/1606.09446>

As the interaction network is likely to contain more than one event, we are interested in finding k events that describe different aspects of the whole network while covering as much activity as possible. This is captured in the following problem formulation.

Problem 3. Maximum k trees (k -MAXTREES): We are given an interaction meta-graph $G = (\mathcal{I}, E, c)$ and $k \in \mathbb{N}$. Find a set of k vertex-disjoint trees $\mathcal{T} = \{T_1, \dots, T_k\}$, with each event tree $T = (V_e, E_e) \in \mathcal{T}$ to be a subgraph of G rooted in some $r_i \in \mathcal{I}$, such that the total number of spanned interactions $|\cup_{T=(V_e, E_e) \in \mathcal{T}} V_e|$ is maximized.

It is easy to observe that this problem is equivalent to maximum k -coverage problem and thus is **NP**-hard. To solve k -MAXTREES efficiently, we consider the question of sampling as few root vertices as possible so that the major events can still be captured. Real-world networks consist of millions of interactions so it is impractical to calculate candidate event trees rooted at each vertex.

5 Algorithms

5.1 Approximating MAXTREE

For finding the best tree, as defined by MAXTREE, we consider three algorithms. Recall that for MAXTREE we are working with the interaction meta-graph G , and that a root vertex is fixed.

Greedy tree growing: The greedy algorithm starts from the root and builds the event tree by adding one vertex (interaction) at a time. At each step the algorithm selects the edge with the minimum cost (topic dissimilarity) from the cutset of the current tree. This choice aims to maximize the topical coherent of the event discovered. The running time is $\mathcal{O}(|\mathcal{I}|^2)$.

Directed Steiner tree algorithm (DST): Recall that MAXTREE corresponds to the *budget* PCST problem. Our second algorithm is inspired by an approach proposed by Johnson et al. [12], where the *budget* PCST problem can be solved by the *quota* PCST problem using binary search. In our case, the prizes of all vertices are uniform, thus the *quota* PCST problem is equivalent to *k*-*minimum spanning tree*. The latter problem can be solved by an algorithm for finding directed Steiner trees (DST), such as the algorithm proposed by Charikar et al. [7]. Thus, our second algorithm uses the DST algorithm, within a binary search to find an event that satisfies the budget constraint. The DST algorithm takes four arguments, G , r , X , and ℓ , where X is a set of terminal nodes and ℓ is a parameter that provides a quality-of-approximation vs. efficiency trade-off. The running time of the algorithm is $\mathcal{O}(|\mathcal{I}|^\ell |X|^{2\ell})$. In our case, $X = \mathcal{I}$, thus the running time is $\mathcal{O}(|\mathcal{I}|^{3\ell})$. We use $\ell = 1$ but still the algorithm is mainly of theoretical interest and not practical for large datasets.

Dynamic programming algorithm (DP): The third algorithm we present is inspired by the idea that when the input DAG is a tree, the problem can be

solved optimally using a simple dynamic programming approach. We investigate two approaches to adapt this algorithm for general (non-tree) DAGs.

In the first approach, we slightly modify the dynamic programming algorithm to make sure the result is a tree. Specifically, when attempting to connect the current node with the subtrees of its children, we enforce the condition that the subtrees cannot have any common nodes. In the second approach, we transform the input DAG into a tree and then apply the original dynamic programming algorithm. Specifically, we first calculate single-source shortest paths from r to all vertices of G using Dijkstra’s algorithm and then apply the dynamic programming algorithm. For integer edge weights and a tree input, the running time is $\mathcal{O}(|\mathcal{I}|B^2)$. In our case, edge weights are real numbers, so we discretize the weights to some decimal digits.

5.2 Approximating k -MAXTREES

Once we have computed a set of candidate event trees using any algorithm for MAXTREE, we need to select k event trees from the candidate set so that vertex coverage is maximized. This is essentially the *maximum coverage* problem. A standard greedy algorithm gives approximation ratio $(1 - \frac{1}{e})$ in time $\mathcal{O}(|\mathcal{I}|^2)$ [17].

5.3 Root sampling strategy

One issue with the greedy max-cover algorithm discussed above, is that all candidate root vertices need to be tested before selecting the one that greedily optimizes the coverage. This is an expensive computational task. To speed up the algorithm for finding top- k trees, we propose a simple root-sampling strategy that ranks roots according to their potential of maximizing MAXTREE.

For every sampled root r we construct a candidate tree D and evaluate *event size upper bound* $U(D, B)$ of DAG D with budget B , defined as:

$$U(D, B) = \max_{F' \in F(D)} \left\{ |F' \cdot \mathcal{I}| \text{ such that } \sum_{e \in F' \cdot E} c(e) \leq B \right\}$$

where $F(D)$ is a set of all forests containing $D.r$ (the root of DAG D).

It is easy to see the optimal tree $T(D, r, B)$ cannot have size greater than $U(D, B)$, thus $U(D, B)$ is indeed an upper bound.

Define the *minimum in-edge* of a vertex u as

$$e^*(G, u) = \arg \min_{e' \in \delta^+(G, u)} c(e'),$$

where $\delta^+(G, u) = \{e \in G.E \mid e.i = u\}$. U can be computed efficiently as follows. Consider only nodes, which belong to $[t_r, t_r + I]$ time interval, where t_r is the root timestamp. Start constructing an event D by adding root r and its child with the lightest edge. Now sort all other nodes by cost of their minimum in-edge cost in increasing order; greedily add nodes with their minimum in-edge to the

Table 1: Network statistics on real datasets. Singleton interactions in the interaction meta-graph are removed.

Datasets	Interaction networks		Interaction meta-graphs		
	#nodes	#edges	#nodes	#edges	Period
<i>Enron</i>	1144	2106	812	21297	1998-10-30 - 2002-02-13
<i>#beefban</i>	11895	33584	26317	75870	2015-03-03 - 2015-03-05
<i>#ukraine</i>	16218	59096	46540	142746	2015-02-27 - 2015-03-03
<i>#baltimore</i>	38541	102139	61501	132012	2015-04-26 - 2015-04-28

event D and stop when budget constraint B is reached. U is a number of nodes in the event D . Note that D is a forest, as we do not care about connectivity during construction.

Our root sampling strategy first ranks all the vertices by U . Then it sequentially selects vertices from the ranked list.

6 Experimental evaluation

As no datasets with ground-truth events are available to us, we validate our approach by using synthetic datasets and by case studies. For the experiments with synthetic datasets: (1) we plant events (considered as ground truth) within random interaction networks; (2) we then apply our algorithm to find events in those synthetic data; (3) we measure the precision and recall of the discovered events with respect to ground-truth. For the case studies we apply our algorithm on *Enron* and on Twitter data, then examine the events we discover, and map them on real known historical events based on textual content and time period. As means of exploratory data analysis, we also visualize the event trees in order to show the information flow within the event.

6.1 Datasets and preprocessing

Synthetic data. We generate synthetic datasets in two steps: (1) we generate ground-truth event trees; (2) we inject noise interactions. Each event is generated independently using the model by Kumar et al. [13], which constructs a tree by iteratively adding random edges. We sample a sender, recipients, timestamp and a topic vector randomly for each node.

Real-world data. We use two real-world datasets: email (*Enron*) and Twitter. Dataset statistics are given in Table 1. **Enron:** we use a preprocessed version of the original Enron dataset [8]. **Twitter:** we use Twitter datasets extracted for three hashtags, each one containing a specific hashtag. The hashtags are *#beefban*, *#baltimore* and *#ukraine*. There is a interaction from a user u to a user v , if the tweet of user v contains username of u . The Twitter datasets are provided by Garimella et al. [9].

Preprocessing. We observe the phenomenon that the same person sends the same (or very similar) messages multiple times, especially on Twitter. Our methods are easily misled by the sheer amount of redundant messages. To avoid this problem, we merge similar messages from the same sender into one. We consider two messages *similar* if (1) they are sent by the same user, (2) their Levenshtein edit distance ratio is below 10%, (3) their time distance is relatively small (e.g., one day). In the newly-merged message, the text content, timestamps are copied from the earliest message. Recipients are the union of all recipients.

We take different approaches for representing interaction content in *Enron* and Twitter. For *Enron*, we train a topic model using *gensim*³. We assign each interaction a topic vector and use cosine distance to compute edge weight.

Measuring tweet similarity is an open challenge due to its short length and conciseness. We took an ensemble approach where vector representation comes from several models. Besides topic vectors, we use also (i) bag-of-word (BoW) with tf-idf re-weighting and (ii) hashtags included in each tweet. For BoW and hashtag representations, we use cosine and Jarccard distance for weight assignment, respectively. Last, we sum up the three distances. For topic modeling, for both *Enron* and Twitter datasets, we use 10 topics, batch size 100 and run it for 10 iterations.

6.2 Results on synthetic datasets

We evaluate five different algorithms for finding the best event: (1) greedy tree growing (*greedy*), (2) binary search using Charikar’s DSP algorithm (*binary_search*), (3) dynamic programming without preprocessing (*DP*), (4) dynamic programming with Dijkstra preprocessing (*DP+dij*), and (5) random tree growing (*random*) as a baseline. The *random* algorithm mimics the *greedy*, but it selects a random edge to grow at each step. We compare quality of solutions obtained on datasets with various noise level. We define *noise level* as a number of noise interactions divided by the total number of interactions of all events. For the DSP algorithm we set level parameter $\ell = 1$, as we have insufficient memory for experiments with larger values.

Different noise levels. To compare the capability of the algorithms to find *one* best event, we generate a sequence of datasets with increasing noise levels and only one event of size 20 (containing 20 nodes). We set ground-truth values of I, B, r for parameters in MAXTREE. We consider three types of measurements: (1) precision, recall, and F1, (2) the value of objective function, and (3) the running time. Log scale is applied in the case of running time as difference between algorithms is of magnitudes order.

In Fig. 2 (a), we see that all our algorithms outperform the trivial *random* baseline. Although *greedy* is a simple heuristic, its performance is among the top. Dijkstra preprocessing for *DP* improves both F1 and computational time. In the contrary, *binary_search* consumes much time, even though it is among the best in other measurements. Notice that *random* achieves high precision because

³ <https://radimrehurek.com/gensim/models/ldamodel.html>




Fig. 2: (a) Performance of the algorithms under noise levels from 0 to 100 with step size 0.5. Results are averaged over 50 repetitions. (b) Performance of the algorithms on synthetic dataset with noise level 20 and varying event size from 10 to 100 at step size 10. Measurement values are averaged from 50 rounds.

it can select a wrong edge that violates the budget constraint at the first few steps and terminate.

Different event sizes. We also study how the algorithms perform in extracting events of different sizes. The experiment setting is similar to the above, but the noise level is fixed to 20, while the event size varies. In Fig. 2 (b), *greedy*, *binary_search* are among the best in terms of precision, recall, F1 and set cover objective, whereas *DP+dij* is slightly worse due to needed edge weight discretization. Again, preprocessing for *DP* improves performance. Running time comparison is consistent with the previous case.

6.3 Parameter effects on real datasets

Effect of B . We evaluate the effect of topic dissimilarity budget B on the tree size objective in MAXTREE. We randomly sample 100 roots for each dataset. B




Fig. 3: Effect of B on the median of tree sizes for different datasets. Note that for $\#ukraine$ and $B > 25$, the DP algorithm fails to complete the experiments as it consumes excessive amount of memory.




Fig. 4: Performance of Different sampling schemes on real datasets: $k = 10$. For Twitter, $B = 15.0, I = 1$ day. For *Enron*, $B = 10.0, I = 4$ weeks. 100 unique roots are selected based on the sampling scheme.

varies from 0 to 100 at a step size of 5.0. For Twitter and *Enron* dataset, I is set to 1 day and 4 weeks respectively. We take the median of all trees returned by each algorithm (Fig. 3).

In *Enron*, we observe a converging effect on both objectives as the dataset is relatively small, while this is not the case in all Twitter datasets. In practice, *greedy* is the best performing algorithm, as it is both competitive in maximizing the objective function and it is computationally efficient.

Sampling scheme comparison. We compare two sampling schemes in real data setting: (1) random root sampling (*random*) as the baseline, (2) ranking roots by event size upperbound (*upperbound*) 5. For each scheme, the set cover objective is recorded whenever a new candidate is added. As we can see in Fig. 4, the event size upper-bound heuristic helps to discover better solutions, especially for *#baltimore* and *Enron*.

Event trees by different algorithms. We compare the behaviours of the algorithms for MAXTREE in real-world datasets. In Fig. 5, the trees are produced by *greedy*, and *DP+dij* are given the same root and budget. The *greedy* algorithm avoids to select heavy edges with weights larger than 0.8 due to its local search strategy whereas *DP+dij* achieves larger tree by selecting a few heavy edges. Therefore we expect *greedy* to produce more topically-coherent events as the pairwise dissimilarity between nodes tend to be smaller.




Fig. 5: Tree computed from *#beefban* given fixed root by *greedy* (a) and *DP+dij* (b), which achieves tree size 46 and 57 respectively. Root, $B = 30$ and $I = 1$ day are the same for both algorithms. Edges with weight ≥ 0.8 are wider. In the tree by *greedy*, no edges with weight ≥ 0.8 are selected. Nodes are colored by senders and edges are colored by its type (broadcast: blue, relay: green, reply: orange)

6.4 Case study in *Enron* dataset

We sample 50 nodes using *upperbound* scheme and applied *greedy* algorithm with $B = 10, I = 28$ days. First, we observed that the events can be grouped into two types: (1) California Energy Crisis,⁴ (2) investigation into Enron’s scandal.⁵ In Fig. 6 (a), we annotated the real world events about the crisis happening during the timespan of the dataset. We found shortly after each major blackout, there is at least one extracted events about it. And before Enron filed bankruptcy, Federal Energy Regulatory Commission (FERC) investigated Enron. Second, in Fig. 7 (a), extracted events tend to occur at the peak of the volume plot.

6.5 Case study in Twitter datasets

We use the same parameters for all three datasets as they have similar size and timespan. Events are extracted by selecting 100 roots using *upperbound* and using *greedy* algorithm with $B = 50$ and $I = 1$ day.

#ukraine. Ukraine crisis arouses media war on Ukraine and Russia.⁶ We observe some of the detected events align well topically and temporally with the actual events in Fig. 6 (b). However, topics are mixed inside some other events. For example, topics on both *#nemstsov* and *#freesavchenko* are detected in event 2. This is expected due to the local similarity measurement in MAXTREE.

#beefban. For the controversial “beef ban”⁷ law in India, results demonstrate clear separation of opinions among events. In Fig. 8, the 1st and 2nd event

⁴ https://en.wikipedia.org/wiki/California_electricity_crisis

⁵ https://en.wikipedia.org/wiki/Enron_scandal

⁶ https://en.wikipedia.org/wiki/Ukrainian_crisis.

⁷ <http://indianexpress.com/article/explained/explained-no-beef-nation/>

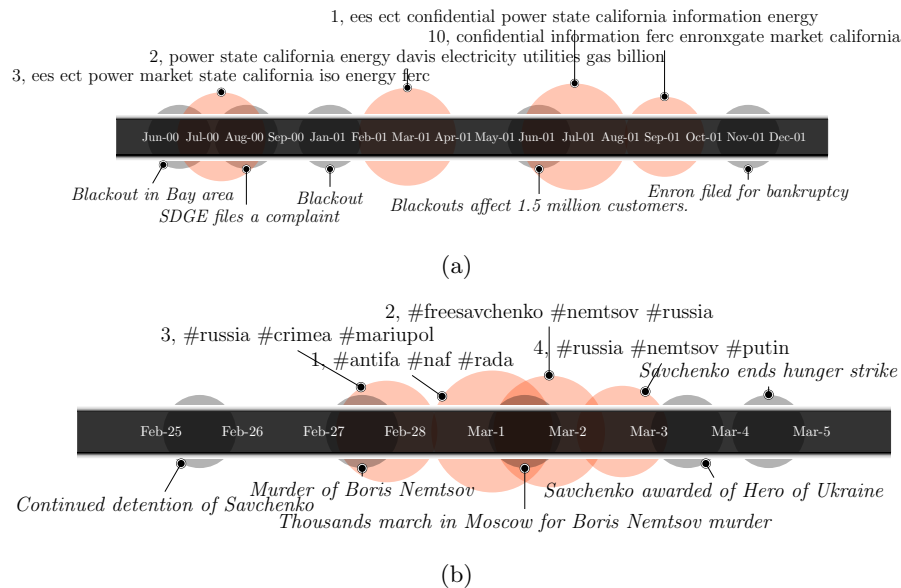


Fig. 6: Timeline with extracted events (larger red circle) and publicly recognised events (smaller black circle and italic text) for Enron (a) and #ukraine (b). (a) highlights events on Enron’s energy scandal and bankruptcy. Event 3, 2, 1 and 10 are displayed. The larger the circle, the larger the event size is. For each event, top topic terms are displayed. In (b), top-4 events are displayed with the top hashtags. Event 2 and 4 maps to the murder of Boris Nemstov (#nemstov), while event 4 also contains tweets on freeing Savchenko (#freesavchenko). Event 1, 3 is about other related issues.

represents opinions opposing and supporting the law. However, we are not able to interpret any temporal pattern in the events due to the short timespan (3 days). We also observe the following. First, certain event (Fig. 8 (a)) display evidence of information propagation. For example, opposing opinions spreads along the user network and affected users also express their objection. Second, for some event (Fig. 8 (b)), dominant user exists who sent more than half of the tweets. Third, we observe events with mixed opinions (Fig. 8 (b)). Last, our method tends to discover events at the “peak” as the set cover objective is better than the “bottom” (Fig. 7 (b)).

#baltimore. We discovered two types of events: (1) “emotional” events showing anger towards the riot, (2) “descriptive” events reporting current situation.

7 Conclusions

We defined the problem of summarizing top- k events in an interaction network. Our approach consists by first transforming the input data into an interaction

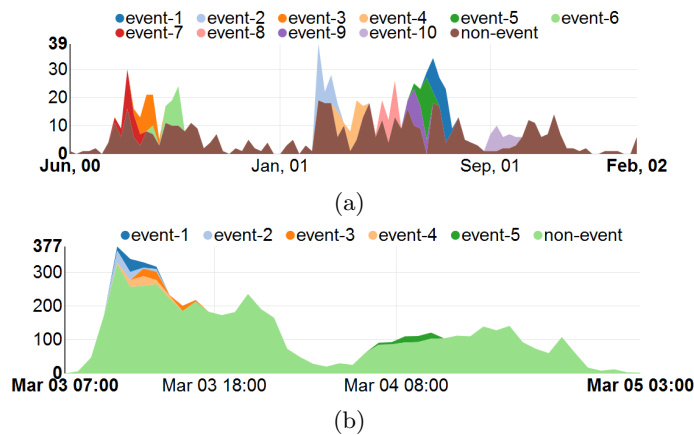


Fig. 7: Stacked area graph of interaction frequency against time. *Enron* (a) contains top-10 events. *#beefban* (b) contains top-5 events.

meta-graph and then defining two optimization problems: budgeted version of PCST and maximum set cover. We offer three algorithms for the former problem. Our experiments show that the greedy approach is more lightweight and performs as good as or even better than other more sophisticated counterparts.

Our work opens many interesting directions for future research. For example, it would be interesting to formulate the problem differently, for example, imposing edge weight constraint for each edge, instead of their weight sum. Another direction is to explore semi-structured interactions such as forums posts, where nesting structures exist between post and comment. We leave scalability experiment and better summarization techniques for future work.

Acknowledgements. This work is partially supported by the Academy of Finland project “Nestor” (286211) and the EC H2020 RIA project “SoBigData” (654024).

References

1. H. Becker, M. Naaman, and L. Gravano. Beyond trending topics: Real-world event identification on twitter. In *ICWSM*, 2011.
2. D. Blei, A. Ng, and M. Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.
3. B. Boden, S. Günnemann, H. Hoffmann, and T. Seidl. Mining coherent subgraphs in multi-layer graphs with edge labels. In *KDD*, 2012.
4. M. Breunig, H.-P. Kriegel, R. Ng, and J. Sander. LOF: identifying density-based local outliers. *SIGMOD record*, 29(2), 2000.
5. M. Cataldi, L. Di Caro, and C. Schifanella. Emerging topic detection on twitter based on temporal and social terms evaluation. In *IWMDM*, 2010.
6. A. Celikyilmaz and D. Hakkani-Tür. Discovery of topically coherent sentences for extractive summarization. In *ACL*, 2011.

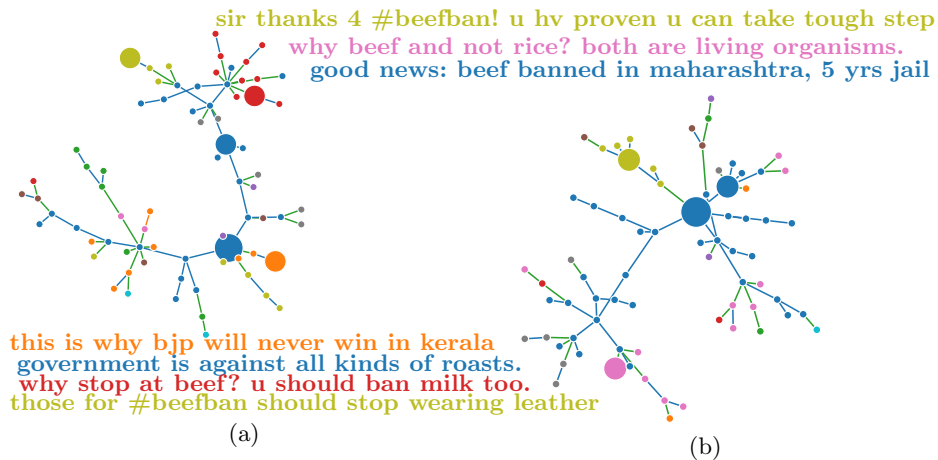


Fig. 8: Extracted events for #beefban. (a) 1st event demonstrates sign of opinion propagation. (b) 2nd event containing mixed opinions. Nodes are colored by the senders. The largest node is the root.

7. M. Charikar, C. Chekuri, T. Cheung, Z. Dai, A. Goel, S. Guha, and M. Li. Approximation algorithms for directed steiner problems. *J. of Algorithms*, 1999.
8. A. Fiore and J. Heer. UC Berkeley Enron email analysis, 2004.
9. K. Garimella, G. De Francisci Morales, A. Gionis, and M. Mathioudakis. Quantifying controversy in social media. In *WSDM*, 2016.
10. V. Guralnik and J. Srivastava. Event detection from time series data. In *KDD*, 1999.
11. D. Inouye and J. K. Kalita. Comparing twitter summarization algorithms for multiple post summaries. In *PASSAT and SocialCom*, 2011.
12. D. Johnson, M. Minkoff, and S. Phillips. The prize collecting Steiner tree problem: theory and practice. In *SODA*, 2000.
13. R. Kumar, M. Mahdian, and M. McGlohon. Dynamics of conversations. In *KDD*, 2010.
14. J. Leskovec, L. Backstrom, and J. Kleinberg. Meme-tracking and the dynamics of the news cycle. In *KDD*, 2009.
15. M. Mathioudakis and N. Koudas. Twittermonitor: trend detection over the twitter stream. In *SIGMOD*, 2010.
16. P. Meladianos, G. Nikolentzos, F. Rousseau, Y. Stavarakas, and M. Vazirgiannis. Degeneracy-based real-time sub-event detection in twitter stream. In *ICWSM*, 2015.
17. V. Vazirani. *Approximation algorithms*. Springer, 2013.
18. J. Weng. Event detection in twitter. In *ICWSM*, 2011.
19. J. Yang and J. Leskovec. Modeling information diffusion in implicit networks. In *ICDM*, 2010.
20. J. Yang and J. Leskovec. Patterns of temporal variation in online media. In *WSDM*, 2011.