

Finding events in temporal networks: Segmentation meets densest-subgraph discovery

Polina Rozenshtein*, Francesco Bonchi[†], Aristides Gionis*, Mauro Sozio[‡] and Nikolaj Tatti[§]

*Aalto University, Espoo, Finland

[†]ISI Foundation, Turin, Italy and Eurecat, Barcelona, Spain

[‡]Telecom ParisTech University, Paris, France

[§]F-Secure Corp., Helsinki, Finland

Abstract—In this paper we study the problem of discovering a timeline of events in a temporal network. We model events as dense subgraphs that occur within intervals of network activity. We formulate the event-discovery task as an optimization problem, where we search for a partition of the network timeline into k non-overlapping intervals, such that the intervals span subgraphs with maximum total density. The output is a sequence of dense subgraphs along with corresponding time intervals, capturing the most interesting events during the network lifetime.

A naïve solution to our optimization problem has polynomial but prohibitively high running time complexity. We adapt existing recent work on dynamic densest-subgraph discovery and approximate dynamic programming to design a fast approximation algorithm. Next, to ensure richer structure, we adjust the problem formulation to encourage coverage of a larger set of nodes. This problem is NP-hard even for static graphs. However, on static graphs a simple greedy algorithm leads to approximate solution due to submodularity. We extended this greedy approach for the case of temporal networks. However, the approximation guarantee does not hold. Nevertheless, according to the experiments, the algorithm finds good quality solutions.

I. INTRODUCTION

Real-world networks are highly dynamic in nature, with new relations (edges) being continuously established among entities (nodes), and old relations being broken. Analyzing the temporal dimension of networks can provide valuable insights about their structure and function, for instance, it can reveal temporal patterns, concept drift, periodicity, temporal events, etc. In this paper we focus on the problem of *finding dense subgraphs*, a fundamental graph-mining primitive. Applications include community detection in social networks [1]–[3], gene expression and drug-interaction analysis in bioinformatics [4], [5], graph compression and summarization [6]–[8], spam and security-threat detection [9], [10], and more.

When working with temporal networks one has first to define how to deal with the temporal dimension, i.e., how to identify which are the temporal intervals in which the dense structures should be sought. Instead of defining those intervals a-priori, in this paper we study the problem of *automatically identifying the intervals that provide the most interesting structures*. We consider a subgraph interesting if it boasts high density. As a result, we are able to discover a sequence of dense subgraphs in the temporal network, capturing the evolution of interesting events that occur during the network lifetime. As a concrete example, consider the problem of *story identification* in online social media [11], [12]: the main goal

is to automatically discover emerging stories by finding dense subgraphs induced by some entities, such as twitter hash-tags, co-occurring in a social media stream. In our case, we are additionally interested in understanding how the stories evolve over time. For instance, as one story wanes and another one emerges, one dense subgraph among entities dissipates and another one appears. Thus, by segmenting the timeline of the temporal network into intervals, and identifying dense subgraphs in each interval, we can capture the evolution and progression of the main stories over time.

As another example, consider a collaboration network, where a sequence of dense subgraphs in the network can reveal information about the main trends and topics over time, along with the corresponding time intervals.

Challenges and contributions. The problem of finding the k densest subgraphs in a static graph has been considered in the literature from different perspectives. One natural idea is to iteratively (and greedily) find and remove the densest subgraphs [13]. More recent works consider finding k densest graphs with limited overlap [14], [15]. However, these approaches do not generalize to temporal networks.

For temporal networks, to our knowledge, there are only few papers that consider the task of finding temporally-coherent densest subgraphs. The most similar to our work aims at finding a heavy subgraph present in all, or k , snapshots [16]. Another related work focuses on finding a dense subgraph covered by k scattered intervals in a temporal network [17]. Both of these methods, however, find a single densest subgraph.

In this paper, instead, we aim at producing a segmentation of the temporal network that (i) captures dense structures in the network; (ii) exhibits temporal cohesion; (iii) spans the whole history of the network; and (iv) is amenable to direct inspection and temporal interpretation. Towards this goal we formulate the problem of k -DENSEST-EPIISODES, which requires to find a partition of the temporal domain into k non-overlapping intervals, such that the intervals span subgraphs with maximum total density. The output is a sequence of dense subgraphs along with corresponding time intervals, capturing the most interesting events during the network lifetime.

A naïve solution to this problem has polynomial but prohibitively-high running-time complexity. Thus, we adapt existing recent work on dynamic-densest subgraph [18] and approximate dynamic programming [19] to design a fast

approximation algorithm.

Next we shift our attention to encouraging coverage of a larger set of nodes, so as to produce richer, more interesting structures. The resulting new problem formulation turns out to be **NP**-hard even for the case of static graphs. However, on static graphs a simple greedy algorithm leads to approximate solution thanks to the submodularity of the objective function. Following this observation, we extended this greedy approach for the case of temporal networks. Despite the fact that the approximation guarantee does not carry on when generalizing to the temporal case, our experimental evaluation indicates that the method produces solutions of very high quality.

The contributions of this paper are summarised as follows:

- We introduce (Section II) the k -DENSEST-EPIISODES problem and show that it has a polynomial time exact algorithm, which is however cubic thus unpractical.
- By leveraging recent work on dynamic densest subgraph and approximate dynamic programming we achieve a fast algorithm with approximation guarantees (Section III).
- We then (Section IV) extend the problem formulation to encourage coverage of a larger set of nodes. We show that the resulting problem is **NP**-hard even for the case of static graph. However, we show on static graphs a simple greedy algorithm leads to approximate solution due to submodularity; then we extend this greedy approach for the case of temporal networks.
- Experiments on synthetic and real-world datasets (Section V), and a case study on Twitter data (Section VI) confirm that our methods are efficient and produce meaningful and high-quality results.

Due to space limitations some of the proofs are omitted and available in the extended version.¹

II. PROBLEM FORMULATION

We are given a *temporal graph* $G = (V, \mathcal{T}, \tau)$, where V denotes the set of nodes, $\mathcal{T} = [0, 1, \dots, t_{max}] \sqsubseteq \mathbb{N}$ is a discrete time domain, and $\tau : V \times V \times \mathcal{T} \rightarrow \{0, 1\}$ is a function defining for each pair of nodes $u, v \in V$ and each timestamp $t \in \mathcal{T}$ whether edge (u, v) exists in t . We denote $E = \{(u, v, t) \mid \tau(u, v, t) = 1\}$ the set of all temporal edges. Given a temporal interval $T = [t_1, t_2] \sqsubseteq \mathcal{T}$, let $G[T] = (V[T], E[T])$ be the subgraph induced by the set of temporal edges $E[T] = \{(u, v) \mid (u, v, t) \in E \wedge t \in T\}$.

Definition 1 (Episode). *Given a temporal graph $G = (V, \mathcal{T}, \tau)$ we define an episode as a pair (I, H) where $I \sqsubseteq \mathcal{T}$ is a temporal interval and H is a subgraph of $G[I]$.*

Our goal is to find a set of interesting episodes along the lifetime of the temporal graph. In particular, our measure of interestingness is the density of the subgraph in the episodes. We adopt the widely-used notion of density of a subgraph $H = (V(H), E(H))$ as the average degree of the nodes in the subgraph, i.e., $d(H) = \frac{|E(H)|}{|V(H)|}$. Observe that this definition is not the only choice, however, such a notion of density enjoys the following nice properties: It can be optimized exactly [20]

and approximated efficiently [21], while a densest subgraph can be computed in real-world graphs containing up to tens of billions of edges [22].

Problem 1 (k -DENSEST-EPIISODES). *Given a temporal graph $G = (V, \mathcal{T}, \tau)$ and an integer $k \in \mathbb{N}$, find a set of k episodes $S = \{(I_\ell, H_\ell)\}$, for $\ell = 1, \dots, k$ such that the $\{I_\ell\}$ are disjoint intervals and $\sum_{\ell=1}^k d(H_\ell)$ is maximized.*

A solution for Problem 1 can be computed in polynomial time. To see this, let S^* be an optimum solution and let $\mathcal{I}(S^*) = \{I_\ell, \ell = 1, \dots, k\}$ and $\mathcal{G}(S^*) = \{H_\ell, \ell = 1, \dots, k\}$. We can assume without loss of generality that the union of the intervals in $\mathcal{I}(S^*)$ results in the set of time stamps \mathcal{T} , that is, $\mathcal{I}(S)$ is a k -segmentation of \mathcal{T} . Moreover, a graph $H_\ell \in \mathcal{G}(S^*)$ is the densest subgraph of $G(I_\ell)$, and can be found in $\mathcal{O}(nm \log n)$ time [20], [23] or in $\mathcal{O}(nm \log(n^2/m))$ time [24] (where n and m denote the number of nodes and edges in $G(I_\ell)$ respectively). The optimal segmentation can be solved with a standard dynamic programming approach, requiring $\mathcal{O}(km^2)$ steps [25]. This brings the total running time to $\mathcal{O}(km^3 n \log n)$ or $\mathcal{O}(km^3 n \log(n^2/m))$.

III. APPROXIMATE DYNAMIC PROGRAMMING

The simple algorithm discussed in the previous section has a running time, which is prohibitively expensive for large graphs. In this section we develop a fast algorithm with approximation guarantees.

The derivations below closely follows the ones in [19], which improves [26]. However, we cannot use those results directly: both papers work with minimization problems, while leveraging the fact that the profit of an interval is not less than the profit of its subintervals (monotone non-decreasing). In contrast, our problem can be viewed as a minimization problem with monotone non-increasing profit function.

Given a time interval $T = [t_1, t_2]$, let us write $d^*(T) = \max_{H \subseteq G(T)} d(H)$. For simplicity, we define $d^*([t_1, t_2]) = 0$ if $t_2 < t_1$. Problem 1 is now a classic k -segmentation problem of \mathcal{T} maximizing the total sum of scores $d^*(T)$ for individual time intervals. For notation simplicity, we assume that the all timestamps \mathcal{T} are enumerated by integers from 1 to r .

Let $o[i, \ell]$ be the profit of optimal ℓ -segmentation using only the first i time stamps. It holds:

$$o[i, \ell] = \max_{j < i} o[j, \ell - 1] + d^*(j + 1, i),$$

and $o[i, k]$ can be computed recursively. Denote the approximate profit of optimal ℓ -segmentation as $s[i, \ell]$. The main idea behind the speed-up is not to test all possible values of j . Instead, we are going to keep a small set of candidates, denoted by A , and only use those values for testing. The challenge is how to keep A small enough while at the same time guarantee the approximation ratio. The pseudo-code achieving this balance is given in Algorithm 1, while a subroutine that keeps the candidate list short is given in Algorithm 2. Algorithm 1 executes a standard dynamic programming search: it assumes that partition of $i' < i$ first data points into $\ell - 1$ intervals is already calculated and finds the best last interval $[a, i]$

¹<https://arxiv.org/abs/1808.09317>

Algorithm 1: APPROXDP(k, ϵ), computes k -segmentation with ϵ -approximation guarantee

Input: number of intervals k , parameter ϵ

Output: approximate solution $s[i, \ell]$ for $i \in [1, r]$,
 $\ell \in [1, k]$

```

1 for  $i = 1, \dots, r$  do  $s[i, 1] = d^*([1, i]);$ 
2 for  $\ell = 2, \dots, k$  do
3    $A = [];$ 
4   for  $i = 1, \dots, r$  do
5     add  $i$  to  $A$ ;
6      $s[i, \ell] = \max\{s[i-1, \ell], s[i, \ell -$ 
7        $1], \max_{a \in A}(s[a-1, \ell-1] + d^*([a, i]))\};$ 
8      $A = \text{SPRS}(A, s[i, \ell], \ell, \epsilon)$ 
9   end
10 return  $s$ 
```

Algorithm 2: SPRS($A, \sigma, \ell, \epsilon$), a subroutine keeping the candidate list short.

Input: candidates A , sparsification factor $\sigma = s[i, \ell]$,
current number of intervals ℓ , approximation
parameter ϵ

Output: sparsified A

```

1  $\delta = \sigma \frac{\epsilon}{k + \ell \epsilon};$ 
2  $j = 1;$ 
3 while  $j < |A| - 1$  do
4   if  $s[a_{j+2}, \ell - 1] - s[a_j, \ell - 1] \leq \delta$  then remove  $a_{j+1}$ 
5     from  $A$ ;
6   else  $j = j + 1;$ 
7 return  $A$ 
```

for partitioning of i first points into ℓ intervals. However, it considers not all possible candidates $[a, i]$, but only a sparsified list, which guarantees to preserve a quality guarantee. The sparsified list is built for a fixed number of intervals ℓ starting from empty list. Intuitively, it keeps only candidates $A = [a_j]$ with significant difference $s[a_j, \ell - 1]$. Significance of the difference depends on the current best profit $s[i, \ell]$: the larger the value of the solution found, the less cautions we can be about lost candidates and the coarser becomes A . Thus, we need to refine A by Algorithm 2 after each processed i .

Let us first prove that APPROXDP yields an approximation guarantee, assuming that $d^*(\cdot)$ is calculated exactly.

Proposition 1. *Let $s[i, \ell]$ be the profit table constructed by APPROXDP(k, ϵ). Then $s[i, \ell](\frac{\ell \epsilon}{k} + 1) \geq o(i, \ell)$.*

To prove the final result, let us first fix ℓ and let A_i be the set of candidates in the beginning of round i . Let δ_i be the value of δ in Algorithm 2, called on iteration i .

Lemma 1. *For every $b \in [1, i - 1]$, there is $a_j, a_{j+1} \in A_i$ with $a_j \leq b \leq a_{j+1}$, such that*

$$s[a_j - 1, \ell - 1] + d^*([a_j, i]) \geq s[b - 1, \ell - 1] + d^*([b, i]) - \delta_{i-1}.$$

Proof. We say that a list of numbers $A = \{a_j\}$ is i -dense, if $s[a_{j+1} - 1, \ell - 1] - s[a_j - 1, \ell - 1] \leq \delta_{i-1}$ or $a_{j+1} = a_j + 1$, (1)

for every $a_j \in A$ with $j < |A|$. We first prove by induction over i that A_i is i -dense.

Assume that A_{i-1} is $(i-1)$ -dense. SPRS never deletes the last element, so $i-1 \in A_{i-1}$, and $A_{i-1} \cup \{i\}$ is $(i-1)$ -dense. Note that $\delta_{i-2} \leq \delta_{i-1}$, because $s[i, \ell]$ is monotonic, $s[i, \ell] \geq s[i-1, \ell]$, due to explicit check on line 5 of APPROXDP. Thus, $A_{i-1} \cup \{i\}$ is i -dense. Since $A_i = \text{SPRS}(A_{i-1} \cup \{i\})$, and SPRS does not create gaps larger than δ_{i-1} , A_i is i -dense.

Let a_j be the largest element in A_i , such that $a_j \leq b$. Then either $a_j \leq b < a_{j+1}$ or $b = a_{|A_i|}$ and $j = |A_i|$. Due to monotonicity, $s[a_{j+1}, \ell - 1] \geq s[b, \ell - 1]$ and gives $s[b - 1, \ell - 1] - s[a_j - 1, \ell - 1] \leq \delta_{i-1}$ for the first case. The second case is trivial.

Due to monotonicity $d^*([a_j, i]) \geq d^*([b, i])$. This concludes the proof. \square

We can now complete the proof.

Proof of Proposition 1. We will prove the result with induction over ℓ . Let $\alpha = (1 + \frac{\epsilon}{k}(\ell - 1))$. Let b be the starting point of the last interval of optimal solution $o[i, \ell]$, and let a_j as given by Lemma 1. Upper bound $\delta_{i-1} = s[i-1, \ell] \frac{\epsilon}{k + \epsilon(\ell - 1)} \leq s[i, \ell] \frac{\epsilon}{k + \epsilon(\ell - 1)} = s[i, \ell] \frac{\epsilon}{\alpha k}$. Then

$$\begin{aligned} \alpha s[i, \ell] &\geq \alpha(s[a_j - 1, \ell - 1] + d^*([a_j, i])) \\ &\geq \alpha(s[b - 1, \ell - 1] + d^*([b, i]) - \delta_{i-1}) \\ &\geq \alpha(s[b - 1, \ell - 1] + d^*([b, i]) - \alpha \delta_{i-1}) \\ &\geq o[i, \ell] - s[i, \ell] \frac{\epsilon}{k}. \end{aligned}$$

As a result, $s[i, \ell](1 + \frac{\epsilon}{k}\ell) \geq o[i, \ell]$. \square

Let us now address the computational complexity.

Proposition 2. *The running time of APPROXDP is $\mathcal{O}(\frac{k^2}{\epsilon}r)$.*

Proof. Fix i and ℓ , and let $c_j = s[a_j, \ell]$, where $a_j \in A_i$. Then c_j is monotonically increasing sequence upper bounded by $s[i-1, \ell]$, and having consecutive elements being at least δ_{i-1} apart. Counting conservatively, this leads to

$$|A_i| \leq 2 + \left\lceil \frac{s[i-1, \ell]}{\delta_{i-1}} \right\rceil \leq 2 + \left\lceil \frac{k(1 + \epsilon)}{\epsilon} \right\rceil \in \mathcal{O}(k/\epsilon).$$

Since we have kr cells in s , the result follows. \square

Since computing d^* requires $\mathcal{O}(nm \log n)$ time, this gives us a total running time of $\mathcal{O}(nmr \frac{k^2}{\epsilon})$. We further speed up our algorithm by approximating the value d^* by means of one of the approaches developed in [18]. In particular, we employ the algorithm that maintains a $2(1 + \epsilon)$ -approximate solution for the incremental densest subgraph problem (i.e. edge insertions only), while boasting a poly-logarithmic amortized cost. We shall refer to such an algorithm as ApprDens.

ApprDens allows us to efficiently maintain the approximate density of the densest subgraph $d^*([a, i])$ for each a in A_i in APPROXDP, as larger values of i are processed and

edges are added. Whenever we remove an item a from A_i in SPDS we also drop the corresponding instance of `ApprDens`.

From the fact that an approximate densest subgraph can be maintained with poly-logarithmic amortized cost, it follows that our algorithm boasts quasi-linear running time.

Proposition 3. *ApproxDP combined with ApprDens runs in $\mathcal{O}(\frac{k^2}{\epsilon_1 \epsilon_2^2} m \log^2 n)$, where ϵ_1 and ϵ_2 are the respective approximation parameters for ApproxDP and ApprDens.*

Proof. Let m_j be the number of edges added to the graph corresponding to a_j before it is deleted. The same argument as in the proof of Proposition 2 states that $\sum_i m_i \in \mathcal{O}(\frac{k^2}{\epsilon_1} m)$. Theorem 4 in [18] states that maintaining the graph with m_i edges requires $\mathcal{O}(m_i \epsilon_2^{-2} \log^2 n)$ time. Combining these two results proves the proposition. \square

When combining `ApproxDP` with `ApprDens`, we wish to maintain the same approximation guarantee of `ApprDens`. Recall that `ApproxDP` leverages the fact that the profit function is monotone non-increasing. Unfortunately, `ApprDens` does not necessarily yield a monotone score function, as the density of the computed subgraph might decrease when a new edge is inserted. This can be easily circumvented by keeping track of the best solution, i.e. the subgraph with highest density. The following proposition holds.

Proposition 4. *ApproxDP combined with ApprDens yields a $2(1 + \epsilon_1)(1 + \epsilon_2)$ -approximation guarantee.*

Proof. Let $d_a^*(T)$ be the density of the graph returned by `ApprDens` for a time interval T . Let O be the optimal k -segmentation, and let $q_1 = \sum_{I \in \mathcal{I}(O)} d^*(O)$ be its score. Let also $q_2 = \sum_{I \in \mathcal{I}(O)} d_a^*(O)$. Let q_3 be the score of the optimal k -segmentation using d_a^* , and let q_4 be the score of the segmentation produced by `ApproxDP`. Then,

$$q_1 \leq 2(1 + \epsilon_2)q_2 \leq 2(1 + \epsilon_2)q_3 \leq 2(1 + \epsilon_2)(1 + \epsilon_1)q_4,$$

completing the proof. \square

We will refer to this combination of `ApproxDP` with `ApprDens` as Algorithm `KGAPPROX`.

IV. ENCOURAGE COVERAGE

Problem 1 is focused on total density maximization, thus its solution can contain graphs which are dense, but union of their node sets cover only a small part of the network. Such segmentation is useful when we are interested in the densest temporally coherent subgraphs which can be understood as tight cores of temporal clusters. However, segmentations with larger but less dense subgraphs, covering a larger fraction of nodes, can be useful to get a high-level explanation of the whole temporal network. To allow for such segmentations we extend Problem 1 to take node coverage into account.

Let $x_v(\mathcal{G}) = |\{G_i \in \mathcal{G} : v \in G_i, G_i \in \mathcal{G}\}|$ be the number of subgraphs in \mathcal{G} , which include node v .

Here we consider a generalized cover functions of the shape

$$\text{cover}(\mathcal{G}) = \sum_{v \in V} w(x_v(\mathcal{G})),$$

where w is a non-negative non-decreasing concave function of $x_v(\mathcal{G})$. When $w(x_v(\mathcal{G}))$ is a 0-1 indicator function, function $\text{cover}(\mathcal{G})$ is a standard cover, which is intuitive and easy to optimize by greedy algorithm. Another instance of the generalized cover function, inspired by text-summarization research [27], is $w(x_v(\mathcal{G})) = \sqrt{x_v(\mathcal{G})}$. It ensures that the marginal gain of a node decreases proportionally to the number of times the node is covered.

Problem 2. *Given a temporal graph $G = (V, E)$, integer k , parameter $\lambda \geq 0$. Find a k -segmentation $S = \{(I_i, G_i)\}$ of G , such that $\text{profit}(S) = \sum_{G_i \in \mathcal{G}} d(G_i) + \lambda \times \text{cover}(\mathcal{G})$ is maximized.*

Proposition 5. *There is no polynomial solution for Problem 2 unless $P=NP$.*

Proposition 6. *Function $\text{cover}(\mathcal{G})$ is a non-negative non-decreasing submodular function of subgraphs.*

Proof. For a fixed $v \in V$ function $x_v(\mathcal{G})$ is non-decreasing modular (and submodular): for any set of subgraphs X and a new subgraph x holds that $x_v(X \cup \{x\}) - x_v(X) = 1$ if v belongs to x and does not belong to any subgraph in X . Otherwise 0. By property of submodular functions, composition of concave non-decreasing and submodular non-decreasing is non-decreasing submodular. Function $\text{cover}(\mathcal{G})$ is submodular non-decreasing as a non-negative linear combination. Non-negativity follows from non-negativity of w . \square

A. *K static densest subgraphs and generalized average degree*

Before going into the temporal segmentation, we briefly consider the static case:

Problem 3. *Given a static graph $H = (V, E')$, integer k , $\lambda \geq 0$. Find a set of k subgraphs $\mathcal{H} = \{H_i \in H\}$, such that $\text{profit}_{ST} = \sum_{H_i \in \mathcal{H}} d(H_i) + \lambda \cdot \text{cover}(\mathcal{H})$ is maximized.*

To solve this problem we can search greedily over subgraphs. Let $\mathcal{H}_{i-1} = \{H_1, \dots, H_{i-1}\}$, and define marginal node gain, given weight function w , as $\delta_v(H_i, \mathcal{H}_{i-1} | w) = w(x_v(\mathcal{H}_{i-1} \cup \{H_i\})) - w(x_v(\mathcal{H}_{i-1}))$. Then denote marginal gain of subgraph H_i given already selected graphs \mathcal{H}_{i-1} as

$$\chi(H_i, \mathcal{H}_{i-1} | w) = d(H_i) + \lambda \sum_{v \in H_i} \delta_v(H_i, \mathcal{H}_{i-1} | w).$$

Greedy algorithm for Problem 3 consequently builds the set \mathcal{H} by adding H_i , which maximizes gain $\chi(H_i, \mathcal{H}_{i-1})$. If we can find H_i optimally, such greedy gives $1-1/e$ approximation due to submodular maximization over cardinality constraints (see [28] for this classic result, Euler's number $e \approx 2.71828$).

Problem 4. *Given a static graph $H = (V, E')$, a set of subgraphs $\mathcal{H}_{i-1} = \{H_1, \dots, H_{i-1}\}$, find graph $H_i \in H$, such that $\chi(H_i, \mathcal{H}_{i-1})$ is maximized.*

Before we proceed, we define a more general and simple version of Problem 4. First, we note that preselected subgraphs \mathcal{H}_{i-1} contribute only to $\delta_v(H_i, \mathcal{H}_{i-1} | w)$ and this term does not change through iterations. Thus, once term $\delta_v(H_i, \mathcal{H}_{i-1} | w)$ is recalculated we can exclude \mathcal{H}_{i-1} from consideration.

Algorithm 3: StaticGreedy

Input: static graph $H = (V, E')$ **Output:** subgraph $\bar{H} \subseteq H$ which maximizes $d_a(\bar{H})$

```

1  $\bar{H} = H;$ 
2 while  $H \neq \emptyset$  do
3    $v = \arg \min_{v \in H} \deg_a(v | H);$ 
4    $H = H \setminus \{v\};$ 
5   if  $d_a(H) > d_a(\bar{H})$  then  $\bar{H} = H;$ 
6 end
7 return  $\bar{H}$ 

```

Next, we define a generalized degree as a function of nodes defined as $\deg_a(v | H) = \sum_{u \in V \setminus \{v\}} a(v, u | H)$ with $a(v, u | H) \geq 0$. Let $\mathbb{I}(v, u | H)$ be 1, if there is an edge between u and v in graph H and 0 otherwise. If $a(v, u | H) = \mathbb{I}(v, u | H)$, then $\deg_a(v | H) = \deg(v | H)$ —a degree of node v in H . Denote a half of the average generalized degree of graph H as $d_a(H) = \frac{1}{2|V(H)|} \sum_{v \in V(H)} \deg_a(v | H)$.

Problem 5. Given a static graph $H = (V, E')$ find graph $H_i = (V(H_i), E'(H_i)) \subseteq H$, such that a half of the average generalized degree $d_a(H_i)$ is maximized.

If $a(v, u | H) = \mathbb{I}(v, u | H)$, then the profit of Problem 5 is the half of average degree $d_a(H_i) = d(H_i)$. On the other hand, when $a(v, u | H) = \mathbb{I}(v, u | H) + 2\lambda|V(H)|\delta_v(H_i, \mathcal{H}_{i-1} | w)$, then Problem 5 is equivalent to Problem 4. Note that in the latter case a depends on the number of nodes in graph H .

We will continue analysis with Problem 5.

Proposition 7. There is no polynomial solution for Problem 5 unless $P=NP$.

To solve Problem 5 efficiently we can modify Charikar’s algorithm for densest subgraphs [21] and obtain 1/2 approximation guarantee.

Proposition 8. Algorithm 3 gives 1/2 approximation for Problem 5.

Time complexity of Algorithm 3 is quadratic of number of nodes (not linear of the edges, like in the case of densest subgraph), as $|V(H)|$ decreases on each step and we need to update generalized degree of all nodes, not only neighbors of the removed node. With a Fibonacci heap time complexity is $\Theta(|V| + |V|(\log |V| + |V|)) = \Theta(|V|^2)$.

B. Incremental case

Here we consider the setting of incremental updates for Problem 5, which may be not interesting by itself, but we will use it as a subroutine for temporal case.

Given a stream of incremental edge updates to graph H we would like to find and keep up-to-date a subgraph H_i , which maximizes $d_a(H_i)$ for some generalized degree function $\deg_a(u, v | H_i)$.

To keep H_i updated we can use the data structure and update procedure designed for the densest subgraph by Epasto et al. [18]. In the full version of this paper we describe the approach of Epasto et al. and necessary modifications to

handle generalized degree. We will refer to this extension as ApprGenDens. The algorithm provides $2(1+\epsilon)$ -approximate generalized density densest subgraph using edge insertions.

Similar to the original algorithm, the generalization requires $O(|V| + |E|)$ of space, while running time increases: $O(\frac{|V|^2}{|E|}\epsilon^{-2} \log^2 D)$ amortized cost per edge insertion, with $D = O(|V|)$ is the maximum of average generalized degree.

C. Greedy dynamic programming

Similarly to Problem 1, we will use dynamic programming for Problem 2. However, as the problem is hard we have to rely on greedy choices of the subgraphs. Thus, the obtained solution does not have any quality guarantee.

Let $M[\ell, i]$ be the profit of i first points into ℓ intervals, let $C[\ell, i]$ be the set of subgraphs $\mathcal{G}_\ell = \{G_1, \dots, G_\ell\}$ selected on these ℓ intervals, $1 \leq \ell \leq k$ and $0 \leq i \leq m$.

Define marginal gain interval $[j, i]$, given that $j - 1$ are already segmented into $\ell - 1$ interval,

$$\text{gain}([j, i], C[\ell - 1, j - 1]) = \max_{G' \subseteq G([j, i])} \chi(G', C[\ell - 1, j - 1]).$$

Dynamic programming recurrence:

$$\begin{aligned}
M[\ell, i] &= \max_{1 \leq j \leq i+1} M[\ell - 1, j - 1] \\
&\quad + \text{gain}([j, i], C[\ell - 1, j - 1]) \text{ for } 1 < \ell \leq k, \\
M[1, i] &= d^*([0, i]) \text{ for } 0 \leq i \leq m, \\
M[k', 0] &= 0 \text{ for } 1 \leq k' \leq k.
\end{aligned}$$

After filling this table, $M[k, m]$ contains the profit of k -segmentation with subgraph overlaps. $C[k, m]$ will contain selected subgraphs, the intervals and subgraphs can be reconstructed, if we keep track of the starting points of selected last intervals. Note, that profit $M[k, m]$ is not optimal, because the choice of subgraph G_i depends on the interval and the previous choices, and there is a fixed order, in which we explore intervals.

We perform dynamic programming by approximation algorithm APPROXDP, and the densest subgraph for each candidate interval is retrieved by ApprGenDens. We refer to the resulting algorithm as KGCVR.

To keep track on number of x_v when we construct \mathcal{G} we need to keep frequencies of each node. To avoid extensive memory costs, in the experiments we use Min-Count sketches.

V. EXPERIMENTS

We evaluate the performance of the proposed algorithms on synthetic graphs and real-world social networks. The datasets are described below. Unless specified, we post-process the output of all algorithms and report the optimal densest subgraphs in the output intervals. Our datasets and implementations are publicly available.²

²<https://github.com/polinalolina/segmentation-meets-densest-subgraph>

A. Synthetic data.

We generate a temporal network with k planted communities and a background network. All graphs are Erdős-Rényi. The communities G' have the same density, disjoint set of nodes, and are planted in non-overlapping intervals. The background network G includes nodes from all planted communities G' . The edges of G are generated uniformly on the timeline. In the typical setup the length of the whole time interval T is $|T| = 1000$ time units, while the edges of each G' are generated in intervals of length $|T'| = 100$ time units. The densities of the communities and the background network vary. The number of nodes in G is set to 100.

We test the ability of our algorithms to discover planted communities in two settings. In the first setting (dataset family *Synthetic1*) we vary the average degree of the background network from 1 to 6 and fix the density of the planted 5-cliques to 4. *Synthetic1* allows to test the robustness against background noise. In the second setting (dataset family *Synthetic2*) we vary the density of planted 8-node graphs from 2 to 7, while the average degree of the background network is fixed to 2.

B. Real-world data.

We use the following real-world datasets: *Facebook* [29] is a subset of Facebook activity in the New Orleans regional community. Interactions are posts of users on each other walls. The data covers the time period from 9.05.06 to 20.08.06. The *Twitter* dataset tracks activity of Twitter users in Helsinki in year 2013. As interactions we consider tweets that contain mentions of other users. The *Students*³ dataset logs activity in a student online network at the University of California, Irvine. Nodes represent students and edges represent messages with ignored directions. *Enron*⁴ is a popular dataset that contains email communication of senior management in a large company and spans several years.

For a case study we create a hashtag network from Twitter dataset (the same tweets from users in Helsinki in year 2013): nodes represent hashtags – there is an interaction, if two hashtags occur in the same tweet. The timestamp of the interaction corresponds to the timestamp of the tweet. We denote this dataset as *Twitter#*.

C. Optimal baseline

A natural baseline for KGAPPROX is OPTIMAL, which combines exact dynamic programming with finding the optimal densest subgraph for each candidate interval. Due to the high time complexity of OPTIMAL we generate a very small dataset with 60 timestamps, where each timestamp contains a random graph with 3–6 nodes and random density. We vary the number of intervals k and report the value of the solution (without any post-processing) and the running time in Figure 1. On this toy dataset KGAPPROX is able to find near-optimal solution, while it is significantly faster than OPTIMAL.

³http://toreopsahl.com/datasets/#online_social_network

⁴<http://www.cs.cmu.edu/~.enron/>

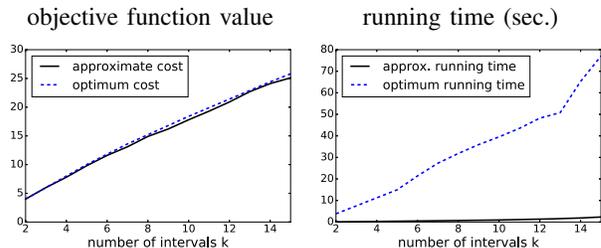


Fig. 1. Comparison between optimum and approximate solutions (OPTIMAL and KGAPPROX). Approximate algorithm was run with $\epsilon_1 = \epsilon_2 = 0.1$. Running time is in seconds.

(a) effect of background noise (b) effect of community density

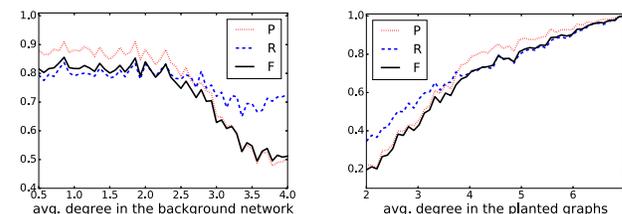


Fig. 2. Precision, recall and F -measure on synthetic datasets. For plot (a) the community average degree is fixed to 5 (*Synthetic1* dataset), for plot (b) the background network degree is fixed to 2 (*Synthetic2* dataset)

D. Results on synthetic datasets

Next, we evaluate the performance of KGAPPROX on the synthetic datasets *Synthetic1* and *Synthetic2* by assessing how well the algorithm finds the planted subgraphs. We report mean precision, recall, and F -measure, calculated with respect to the ground-truth subgraphs. All results are averaged over 100 independent runs.

First, Figure 2(a) depicts the quality of the solution as a function of background noise. Recall, that the *Synthetic1* dataset contains planted 8-node subgraphs with average degree 5. Precision and recall are generally high for all values of average degree in the background network. However, precision degrades as the density of the background network increases, as then it becomes cost-beneficial to add more nodes in the discovered densest subgraphs.

Second, Figure 2(b) shows the quality of the solution of KGAPPROX as a function of the density in the planted subgraphs. In *Synthetic2* the density of the background is 2. Similarly to the previous results, the quality of the solution, especially recall, degrades much only when the density of the planted and the background network become similar.

E. Results on real-world datasets

As the optimal partition algorithm OPTIMAL is not scalable for real datasets, we present comparative results of KGAPPROX with baselines KGOPTDP and KGOPTDS. The KGOPTDP algorithm performs exact dynamic programming, but uses an approximate incremental algorithm for the densest subgraph search (the incremental framework by Epasto et al. [18]). Vice versa, KGOPTDS performs approximate dynamic programming while calculating the densest subgraph optimally for

each candidate interval (by Goldberg’s algorithm [20]). Note that KGOPTDP has $2(1 + \epsilon_{DS})^2$ approximation guarantee and KGOPTDS has $(1 + \epsilon_{DP})$ approximation guarantee. However, even these non-optimal baselines are quite slow on practice and we use a subset of 1000 interactions of *Students* and *Enron* datasets for comparative reporting.

To ensure fairness, we report the total density of the optimal densest subgraphs in the intervals returned by the algorithms.

In Table I we report the density of the solutions reported by KGAPPROX, KGOPTDP, and KGOPTDS, as well as their running time. We experiment with different parameters for the approximate densest-subgraph search (ϵ_{DS}) and for approximate dynamic programming (ϵ_{DP}).

For both datasets the best solution was found by KGOPTDS. This is expected, as this algorithm has the best approximation factor. The solution cost decreases as ϵ_{DP} increases. On the other hand, KGOPTDS has the largest running time, which decreases with increasing ϵ_{DP} , but even with the largest parameter value ($\epsilon_{DP} = 2$) KGOPTDS takes about an hour.

The KGOPTDP algorithm typically finds the second-best solution, however it only marginally outperforms KGAPPROX (e.g., $\epsilon_{DS} = 0.1$), while requiring up to several orders of magnitude of higher computational time. Naturally, the quality of the solution degrades with increasing ϵ_{DS} .

The solution quality degrades with increasing the approximation parameters for all algorithms. However, the degradation is not as dramatic as the worst case bound suggests, and using such an approximation parameter offers significant speed-up. KGAPPROX provides the fastest estimates of a good quality for a wide range of approximation parameters. Note that KGAPPROX is more sensitive to the changes in the quality of the densest subgraph search regulated by ϵ_{DS} .

F. Running time and scalability

Figure 3 shows running time of KGAPPROX as a function of the approximation parameters ϵ_{DS} and ϵ_{DP} . The figure confirms the theory, that is, ϵ_{DS} has significant impact on the running time, while the algorithm scales very well with ϵ_{DP} .

We demonstrate scalability in Figure 4, plotting the running time for increasing number of interactions, for *Facebook* and *Twitter* datasets. Recall that the theoretical running time is $\mathcal{O}(k^2 m \log n)$, where n is the number of nodes and m the number of interactions. In practice, the running time grows fast for the first thousand interactions and then saturates to linear dependence. This happens because in the beginning of the network history the number of nodes grows fast. In addition, new, denser than previously seen, subgraphs are more likely to occur. Thus, the approximate densest-subgraph subroutine has to be computed more often. Furthermore, the number of intervals k contributes to running time as expected.

G. Subgraphs with larger node coverage — static graphs

Next we evaluate STATICGREEDY. To measure coverage, we simply count the number of distinct nodes in the output subgraphs. We use the 10K first interactions of *Students* dataset, set $k = 20$, and test different values of λ . Figure 5 shows the density and the pairwise Jaccard similarity of the

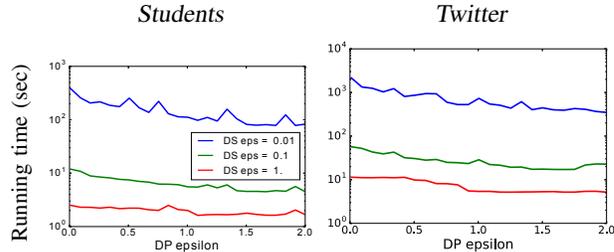


Fig. 3. Effect of different approximation parameters in KGAPPROX. $k = 20$.

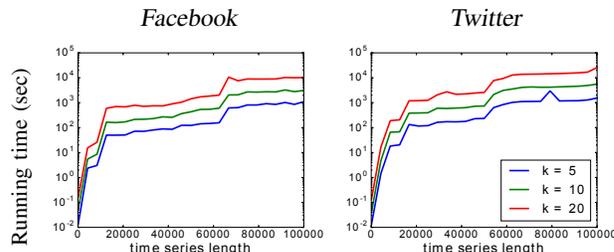


Fig. 4. Scalability testing with $\epsilon_{DS} = \epsilon_{DP} = 0.1$.

node sets of the retrieved subgraphs. The subgraphs are shown in the order they are discovered. Smaller values of λ give larger density, and larger values of λ give more cover. We observe that, for all values of λ , in the beginning STATICGREEDY returns diverse and dense subgraphs, but soon after it returns identical graphs. We speculate that the algorithm finds all dense subgraphs that exist in the dataset. Regarding setting λ , we observe that $\lambda = 0.002$ offers a good trade-off in finding subgraphs of high density and moderate overlap.

H. Subgraphs with larger node coverage — dynamic graphs

Finally we evaluate the performance of KGCVR algorithm. We vary the parameter λ and compare different characteristics of the solution, with the solution returned by KGAPPROX. For different values of λ , Table II shows average density, total number of covered nodes, average size of the subgraphs, and average pairwise Jaccard similarity. Although KGCVR does not have an approximation guarantee, for small values of λ it finds subgraphs of the density close to KGAPPROX. Similarly to the static case, λ provides an efficient trade-off between density and coverage.

VI. CASE STUDY

We present a case study using graphs of co-occurring hashtags from Twitter messages in the Helsinki region. We create two subsets of *Twitter#* dataset: one covering all tweets in November 2013 and another in December 2013. Figure 6 shows the dense subgraphs discovered by the KGAPPROX algorithm on these datasets, with $k = 4$ and $\epsilon_{DS} = \epsilon_{DP} = 0.1$.

For the November dataset, KGAPPROX creates a small 1-day interval in the beginning and then splits the rest time almost evenly. This first interval includes the nodes `movember`, `liiga`, `halloween`, and `digixpo`, which cover a broad range of global (e.g., `movember` and `Halloween`) and local

TABLE I
COMPARISON WITH KGOPDP AND KGOPDTS BASELINES.

Dataset		Community density					Running time (sec.)						
		KGAPPROX	ϵ_{DS}				KGOPDTS	KGAPPROX	ϵ_{DS}				KGOPDTS
		0.01	0.1	1	2		0.01	0.1	1	2			
Students 1000	ϵ_{DP}	0.01	4.24	4.24	3.82	3.82	6.30	0.01	0.62	0.62	0.63	0.64	23678
		0.1	4.24	4.24	3.82	3.82	6.22	0.1	0.23	0.23	0.24	0.23	11877
		1	4.24	4.24	3.82	3.82	5.76	1	0.13	0.26	0.13	0.13	3394
		2	4.24	4.24	3.82	3.82	5.61	2	0.36	0.20	0.20	0.36	3769
		KGOPDP	5.73	5.73	3.82	3.82		KGOPDTS	162	43.5	29.5	29.5	
	KGOPDTS												
Enron 1000	ϵ_{DP}	0.01	10.4	10.4	10.0	10.5	11.3	0.01	56.4	55.5	42.3	31.8	25788
		0.1	10.3	10.4	10.0	10.3	11.0	0.1	3.02	2.85	2.07	1.70	16070
		1	9.54	9.54	8.80	9.83	11.0	1	0.43	0.44	0.29	0.28	7834
		2	7.34	7.34	7.34	7.34	10.8	2	0.22	0.22	0.23	0.23	3469
		KGOPDP	10.5	11.0	10.4	10.4		KGOPDTS	1654	61.15	17.82	6.07	
	KGOPDTS												

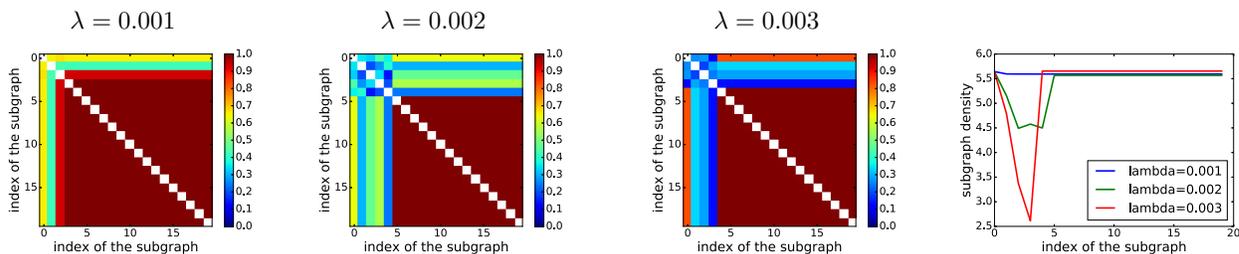


Fig. 5. Pairwise similarities (3 heatmap plots on the left) and densities (right plot) of subgraphs returned by STATICGREEDY.

events (e.g., game-industry event DigiExpo and Finnish ice-hockey league). The next interval is represented by a large variety of well-connected tags related to `mtv` and `media`, corresponding to the MTV Europe Music Awards’13 on November 10. There are also other ice hockey-related tags, e.g., `leijonat`, and Father’s Day tags, e.g., `isänpäivä`, which was on November 13. The third interval is mostly represented by Slush-related tags; Slush is the annual large startup and tech event in Helsinki. The last interval is completely dedicated to ice-hockey with many team names.

There are three major public holidays in December: Finland’s Independence Day on December 6, Christmas on December 25, and New Year’s Eve on December 31. KGAPPROX allocates one interval for Christmas and New Year from December 21 to 31. Ice hockey is also represented in this interval, as well as in the third interval. Remarkably, the Independence Day holiday is split into 2 intervals. The first one is from December 1 to December 6, 3:30pm, and the corresponding graph has two clusters: the first one containing general holidays-related tags and the second one is focused on Independence Day President’s reception. This is a large event that starts on December 6, 6pm, is broadcasted live, and is discussed in media for the following days. The second interval for December 6-9 is a truthful representation of this event.

VII. RELATED WORK

Partitioning a graph in dense subgraphs is a well-established problem. Many of the existing works adopt as density definition the average-degree notion [30]–[33]. The densest subgraph, under this definition, can be found in polynomial time [20]. Moreover, there is a 2-approximation greedy algorithm by Charikar [21] and Asahiro [34], which runs in linear time of the graph size. Many recent works develop methods to maintain the average-degree densest-subgraph in a streaming scenario [18], [35]–[38]. Alternative density definitions, such as variants of quasi-clique, are often hard to approximate or solve by efficient heuristics due to connections to **NP**-complete Maximum Clique problem [13], [39], [40].

A line of work focuses on dynamic graphs, which model node/edge additions/deletions. Different aspects of network evolution, including evolution of dense groups, were studied in this setting [41]–[44]. However, here we use the interaction-network model, which is different to dynamic graphs, as it captures the instantaneous interactions between nodes.

Another classic approach to model temporal graphs is to consider graph snapshots, find structures in each snapshot separately (or by incorporating information from previous snapshots), and then summarize historical behavior of the discovered structures [45]–[49]. These approaches usually focus on the temporal coherence of the dense structures discovered

TABLE II
RESULTS OF KGCVR WITH $k = 5$ AND $\epsilon_{DS} = \epsilon_{DP} = 0.1$.

Dataset	λ	Density		Cover		Size		JSim	
		KGCVR	KGAPPROX	KGCVR	KGAPPROX	KGCVR	KGAPPROX	KGCVR	KGAPPROX
Students	1e-6	10.690	11.151	136	130	48.75	37.6	0.1449	0.0951
	1e-5	7.0869	11.151	813	130	261.0	37.6	0.0788	0.095
	1e-4	5.0273	11.151	889	130	286.0	37.6	0.0910	0.0951
Enron	1e-6	19.995	19.871	38	37	16.0	16.2	0.3619	0.3851
	1e-5	19.962	19.871	40	37	17.0	16.2	0.3660	0.3851
	1e-4	6.5684	19.871	1144	37	288.8	16.2	0.0808	0.3851
Facebook	1e-8	5.3714	5.3933	83	120	22.75	27.6	0.0185	0.0163
	1e-7	4.2749	5.3933	3470	120	882.0	27.6	0.0027	0.0163
	1e-6	3.2673	5.3933	4100	120	1228.75	27.6	0.0335	0.0163
Twitter	1e-7	9.9970	10.138	128	152	44.25	54.0	0.1590	0.1673
	1e-6	6.5500	10.138	3808	152	1061.75	54.0	0.0837	0.1673
	1e-5	3.5389	10.138	4604	152	1379.0	54.0	0.0773	0.1673

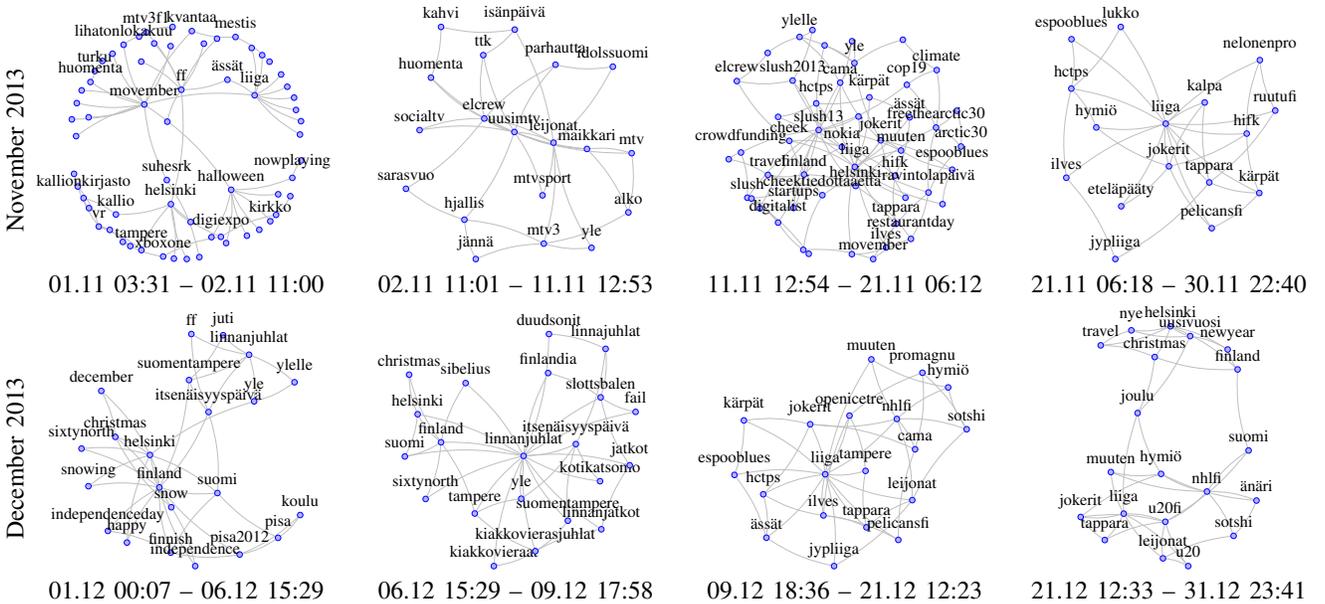


Fig. 6. Subgraphs, discovered in the network of Twitter hashtags $Twitter\#$ by KGAPPROX algorithm with $k = 4$, $\epsilon_{DS} = \epsilon_{DP} = 0.1$.

in the snapshots and assume that the snapshots are given. In this work we aggregate instantaneous interaction into timeline partitions of arbitrary lengths.

To the best of our knowledge, the following works are better aligned with our approach. A work of Rozenshtein et al. [17] considers a problem of finding the densest subgraph in a temporal network. However, first, they do not aim at creating a temporal partitioning. Second, they are interested in finding a single dense subgraph whose edges occur in k short time intervals. On the contrary, in this work we search for an interval partitioning and consider only graphs that are spanned by continuous intervals. Other close works are by Jethava and Beerenwinkel [50] and Semertzidis et al. [16]. However, these works consider a set of snapshots and search for a single

heavy subgraph induced by one or several intervals. The work of Semertzidis et al. [16] explores different formulations for the *persistent heavy subgraph problem*, including maximum average density, while Jethava and Beerenwinkel [50] focus solely on maximum average density.

VIII. CONCLUSIONS

In this work we consider the problem of finding a sequence of dense subgraphs in a temporal network. We search for a partition of the network timeline into k non-overlapping intervals, such that the intervals span subgraphs with maximum total density. To provide a fast solution for this problem we adapt recent work on dynamic densest subgraph and approximate dynamic programming. In order to ensure that the episodes

we discover consist of a diverse set of nodes, we adjust the problem formulation to encourage coverage of a larger set of nodes. While the modified problem is **NP**-hard, we provide a greedy heuristic, which performs well on empirical tests.

The problems of temporal event detection and timeline segmentation can be formulated in various ways depending on the type of structures that are considered to be interesting. Here we propose segmentation with respect to maximizing subgraph density. The intuition is that those dense subgraphs provide a sequence of interesting events that occur in the lifetime of the temporal network. However, other notions of interesting structures, such as frequency of the subgraphs, or statistical non-randomness of the subgraphs, can be considered for future work. In addition, it could be meaningful to allow more than one structure per interval. Another possible extension is to consider overlapping intervals instead of a segmentation.

ACKNOWLEDGMENTS

Part of this work was done while the first author was visiting ISI Foundation. This work was partially supported by three Academy of Finland projects (286211, 313927, and 317085), and the EC H2020 RIA project “SoBigData” (654024). We thank the anonymous reviewers for their valuable comments.

REFERENCES

- [1] J. Chen and Y. Saad, “Dense subgraph extraction with application to community detection,” *TKDE*, vol. 24, no. 7, pp. 1216–1230, 2012.
- [2] D. DiTursi, G. Ghosh, and P. Bogdanov, “Local community detection in dynamic networks,” *arXiv:1709.04033*, 2017.
- [3] D. Taylor, R. S. Caceres, and P. J. Mucha, “Super-resolution community detection for layer-aggregated multilayer networks,” *Physical Review X*, vol. 7, no. 3, p. 031056, 2017.
- [4] E. Fratkin, B. T. Naughton, D. L. Brutlag, and S. Batzoglou, “Motifcut: regulatory motifs finding with maximum density subgraphs,” *Bioinformatics*, vol. 22, no. 14, pp. e150–e157, 2006.
- [5] B. Saha, A. Hoch, S. Khuller, L. Raschid, and X.-N. Zhang, “Dense subgraphs with restrictions and applications to gene annotation graphs,” in *RECOMB*, 2010.
- [6] T. Feder and R. Motwani, “Clique partitions, graph compression and speeding-up algorithms,” *JCSS*, vol. 51, no. 2, pp. 261–272, 1995.
- [7] C. Karande, K. Chellapilla, and R. Andersen, “Speeding up algorithms on compressed web graphs,” *Internet Mathematics*, 2009.
- [8] C. Hernández and G. Navarro, “Compressed representation of web and social networks via dense subgraphs,” in *SIGIR*, 2012, pp. 264–276.
- [9] D. Gibson, R. Kumar, and A. Tomkins, “Discovering large dense subgraphs in massive graphs,” in *PVLDB*, 2005, pp. 721–732.
- [10] A. Beutel, W. Xu, V. Guruswami, C. Palow, and C. Faloutsos, “Copycatch: stopping group attacks by spotting lockstep behavior in social networks,” in *WWW*, 2013, pp. 119–130.
- [11] A. Angel, N. Sarkas, N. Koudas, and D. Srivastava, “Dense subgraph maintenance under streaming edge weight updates for real-time story identification,” *PLVDB*, vol. 5, no. 6, pp. 574–585, 2012.
- [12] O. D. Balalau, C. Castillo, and M. Sozio, “Evidense: A graph-based method for finding unique high-impact events with succinct keyword-based descriptions,” in *Proceedings of the Twelfth International Conference on Web and Social Media, ICWSM*, 2018, pp. 560–563.
- [13] C. Tsourakakis, F. Bonchi, A. Gionis, F. Gullo, and M. Tsiarli, “Denser than the densest subgraph: extracting optimal quasi-cliques with quality guarantees,” in *KDD*, 2013, pp. 104–112.
- [14] O. D. Balalau, F. Bonchi, T. Chan, F. Gullo, and M. Sozio, “Finding subgraphs with maximum total density and limited overlap,” in *WSDM*, 2015, pp. 379–388.
- [15] E. Galbrun, A. Gionis, and N. Tatti, “Top-*k* overlapping densest subgraphs,” *DMKD*, vol. 30, no. 5, pp. 1134–1165, 2016.
- [16] K. Semertzidis, E. Pitoura, E. Terzi, and P. Tsaparas, “Best friends forever (bff): Finding lasting dense subgraphs,” *arXiv:1612.05440*, 2016.
- [17] P. Rozenstein, N. Tatti, and A. Gionis, “Finding dynamic dense subgraphs,” *TKDD*, vol. 11, no. 3, p. 27, 2017.
- [18] A. Epasto, S. Lattanzi, and M. Sozio, “Efficient densest subgraph computation in evolving graphs,” in *WWW*, 2015, pp. 300–310.
- [19] N. Tatti, “Strongly polynomial efficient approximation scheme for segmentation,” *ArXiv:1805.11170*, 2018.
- [20] A. V. Goldberg, *Finding a maximum density subgraph*. University of California Berkeley, 1984.
- [21] M. Charikar, “Greedy approximation algorithms for finding dense components in a graph,” in *APPROX*, 2000, pp. 84–95.
- [22] M. Danisch, T. H. Chan, and M. Sozio, “Large scale density-friendly graph decomposition via convex programming,” in *Proceedings of the 26th International Conference on World Wide Web, WWW 2017*, 2017.
- [23] J. B. Orlin, “Max flows in $o(nm)$ time, or better,” in *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, 2013.
- [24] G. Gallo, M. D. Grigoriadis, and R. E. Tarjan, “A fast parametric maximum flow algorithm and applications,” *SIAM J. Computing*, 1989.
- [25] R. Bellman, *Dynamic programming*. Courier Corporation, 2013.
- [26] S. Guha, N. Koudas, and K. Shim, “Data-streams and histograms,” in *STOC*, 2001, pp. 471–475.
- [27] H. Lin and J. Bilmes, “A class of submodular functions for document summarization,” in *ACL*, 2011, pp. 510–520.
- [28] G. Nemhauser, L. Wolsey, and M. Fisher, “An analysis of approximations for maximizing submodular set functions,” *Mathematical Programming*, vol. 14, no. 1, pp. 265–294, 1978.
- [29] B. Viswanath, A. Mislove, M. Cha, and K. Gummadi, “On the evolution of user interaction in facebook,” in *WOSN*, 2009, pp. 37–42.
- [30] R. Andersen and K. Chellapilla, “Finding dense subgraphs with size bounds,” in *WAW*, 2009, pp. 25–37.
- [31] S. Khuller and B. Saha, “On finding dense subgraphs,” in *ICALP*, 2009.
- [32] E. Galbrun, A. Gionis, and N. Tatti, “Overlapping community detection in labeled graphs,” *DMKD*, vol. 28, no. 5-6, pp. 1586–1610, 2014.
- [33] C. E. Tsourakakis, “A novel approach to finding near-cliques: The triangle-densest subgraph problem,” *arXiv:1405.1477*, 2014.
- [34] Y. Asahiro, K. Iwama, H. Tamaki, and T. Tokuyama, “Greedily finding a dense subgraph,” *Journal of Algorithms*, vol. 34, no. 2, 2000.
- [35] S. Bhattacharya, M. Henzinger, D. Nanongkai, and C. Tsourakakis, “Space- and time-efficient algorithm for maintaining dense subgraphs on one-pass dynamic streams,” in *STOC*, 2015, pp. 173–182.
- [36] H. Esfandiari, M. Hajiaghayi, and D. Woodruff, “Applications of uniform sampling: Densest subgraph and beyond,” *arXiv:1506.04505*, 2015.
- [37] A. McGregor, D. Tench, S. Vorotnikova, and H. T. Vu, “Densest subgraph in dynamic graph streams,” in *MFCSS*. Springer, 2015.
- [38] M. Mitzenmacher, J. Pachocki, R. Peng, C. Tsourakakis, and S. C. Xu, “Scalable large near-clique detection in large-scale networks via sampling,” in *KDD*, 2015, pp. 815–824.
- [39] K. Makino and T. Uno, “New algorithms for enumerating all maximal cliques,” in *SWAT*, 2004, pp. 260–272.
- [40] J. I. Alvarez-Hamelin, L. Dall’Asta, A. Barrat, and A. Vespignani, “Large scale networks fingerprinting and visualization using the *k*-core decomposition,” in *NIPS*, 2006.
- [41] L. Backstrom, D. Huttenlocher, J. Kleinberg, and X. Lan, “Group formation in large social networks: membership, growth, and evolution,” in *KDD*, 2006, pp. 44–54.
- [42] M. Berlingerio, F. Bonchi, B. Bringmann, and A. Gionis, “Mining graph evolution rules,” in *ECML PKDD*, 2009, pp. 115–130.
- [43] S. A. Myers and J. Leskovec, “The bursty dynamics of the twitter information network,” in *WWW*, 2014, pp. 913–924.
- [44] R.-H. Li, J. X. Yu, and R. Mao, “Efficient core maintenance in large dynamic graphs,” *TKDE*, vol. 26, no. 10, pp. 2453–2465, 2014.
- [45] Y.-R. Lin, Y. Chi, S. Zhu, H. Sundaram, and B. L. Tseng, “Facenet: a framework for analyzing communities and their evolutions in dynamic networks,” in *WWW*, 2008, pp. 685–694.
- [46] S. Asur, S. Parthasarathy, and D. Ucar, “An event-based framework for characterizing the evolutionary behavior of interaction graphs,” *TKDD*, vol. 3, no. 4, p. 16, 2009.
- [47] D. Greene, D. Doyle, and P. Cunningham, “Tracking the evolution of communities in dynamic social networks,” in *ASONAM*, 2010.
- [48] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela, “Community structure in time-dependent, multiscale, and multiplex networks,” *Science*, vol. 328, no. 5980, pp. 876–878, 2010.
- [49] M. Berlingerio, F. Pinelli, and F. Calabrese, “Abacus: frequent pattern mining-based community discovery in multidimensional networks,” *DMKD*, vol. 27, no. 3, pp. 294–320, 2013.
- [50] V. Jethava and N. Beerenwinkel, “Finding dense subgraphs in relational graphs,” in *ECML PKDD*, 2015, pp. 641–654.