

AUGSP LICING: Synchronized Behavior Detection in Streaming Tensors

Jiabao Zhang¹, Shenghua Liu^{1#}, Wenting Hou², Siddharth Bhatia³,
Huawei Shen¹, Wenjian Yu^{4#}, Xueqi Cheng¹

¹ Institute of Computing Technology, Chinese Academy of Sciences

² Beijing InnovSharing Co.Ltd

³ National University of Singapore

⁴Dept. Computer Science & Tech., Tsinghua University

{zhangjiabao,liushenghua,shenhuawei,cxq}@ict.ac.cn, siddharth@comp.nus.edu.sg, yu-wj@tsinghua.edu.cn

Abstract

How can we track synchronized behavior in a stream of time-stamped tuples, such as mobile devices installing and uninstalling applications in the lockstep, to boost their ranks in the app store? We model such tuples as entries in a streaming tensor, which augments attribute sizes in its modes over time. Synchronized behavior tends to form dense blocks (i.e. subtensors) in such a tensor, signaling anomalous behavior, or interesting communities. However, existing dense block detection methods are either based on a static tensor, or lack an efficient algorithm in a streaming setting. Therefore, we propose a fast streaming algorithm, AUGSP LICING, which can detect the top dense blocks by incrementally splicing the previous detection with the incoming ones in new tuples, avoiding re-runs over all the history data at every tracking time step. AUGSP LICING is based on a splicing condition that guides the algorithm (Section 4). Compared to the state-of-the-art methods, our method is (1) effective to detect fraudulent behavior in installing data of real-world apps and find a synchronized group of students with interesting features in campus Wi-Fi data; (2) robust with splicing theory for dense block detection; (3) streaming and faster than the existing streaming algorithm, with closely comparable accuracy.

1 Introduction

Given a stream of time-stamped tuples $(a_1, a_2, \dots, a_n, t)$, how can we spot the most synchronized behavior up to now in real-time?

Such a problem has many real applications. In online review sites such as Yelp, let a_1 be a user, a_2 be a restaurant, a_3 be a rating score, and t be the rating time. The most synchronized rating behavior of high scores indicates the most suspicious review fraud (Hooi et al. 2016; Jiang et al. 2014). In application logs, a_1 , a_2 , a_3 , and t can represent a mobile device, an app, installing time, and uninstalling time respectively. Highly synchronous installation and uninstallation from a group of devices can reveal the most suspicious behavior of boosting target apps' ranks in an app store. In terms of pattern discovery, synchronous connections and disconnections to the Wi-Fi access point (AP) in campus Wi-Fi connection logs can discover students that have the same classes of interest.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Corresponding authors

Let such time-stamped tuples be entries of a tensor with multiple dimensions, such as *user*, *object*, and *time* (Figure 1a). Note that we call each dimension as a mode like (Lu, Plataniotis, and Venetsanopoulos 2013), and a two-mode tensor is a matrix. Since tensors allow us to consider additional information especially the time, the densest block (subtensor) of interest can identify the most synchronized behavior in time-stamped tuples (Jiang et al. 2015; Shah et al. 2015; Shin et al. 2017a).

In such a streaming tensor, the attribute size of *time* mode is augmented over time as shown in Figure 1a. Other modes such as *user* and *object* can also be augmented when an unseen user or object is observed. Nowadays, dense subtensor detection methods for streaming tensors are essential. This is because it is much easier than in the past to collect large datasets with the advance of technology. Not only is the size of real data very large, but also the rate at which it arrives is high (Akoglu, Tong, and Koutra 2015). For example, Facebook users generate billions of posts every day, billions of credit card transactions are performed each day, and so on. As such, whole data may be too large to fit in memory or even on a disk. On the other hand, we can think of this kind of data generation as streaming tensors as mentioned above. Thus, the methods which can update their estimations efficiently when the tensor changes over time are essential for dense subtensor detection problem. However, many existing works on dense subtensor detection were designed for static tensors given in a batch (Shin, Hooi, and Faloutsos 2016; Shin et al. 2017a; Yikun et al. 2019) and we refer to them as batch algorithms. Although these batch algorithms are near-linear with the size of tuples (i.e. non-zero entries) in a tensor, re-running the algorithms at every time step for a streaming tensor can result in memory overload when we meet huge size datasets and quadratic time complexity. This causes limited scalability in a streaming setting due to the repeated computation on past tuples (Teng 2016). As for the state-of-the-art streaming algorithm, DENSESTREAM (Shin et al. 2017b), maintained a fine-grained order (i.e. D-order) to search for the densest subtensor. The order is updated for every single new tuple, limiting the detection speed.

Therefore we propose AUGSP LICING, a fast and incremental algorithm to approximate the up-to-date dense blocks in streaming tensors. Without re-running batch algorithms, our heuristic algorithm based on the splicing condition re-

duces the search space, incrementally splices dense blocks of previous detections and the new blocks detected only in an incoming tensor (right-side tensor in Figure 1a). As such, AUGSPICING can detect dense subtensors at every time step in real-time. Experiments show that AUGSPICING is the fastest, with comparable accuracy with the state-of-the-art methods. In summary, our main contributions are:

1. **Fast and Streaming Algorithm:** We propose a fast dense block detection algorithm in streaming tensors, which is up to 320 times faster than the current state-of-the-art algorithms (Figure 1b).
2. **Robustness:** AUGSPICING is robust with splicing theory to do incremental splices for dense block detection.
3. **Effectiveness and Explainable Detection:** Our algorithm achieves accuracy (in terms of F-measure) comparable to the best baseline, DENSESTREAM (Figure 1b). AUGSPICING spots suspicious mobile devices that boost target apps' ranks in a recommendation list by synchronous installations and uninstalls in real-world data. The result shows that the suspicious installations of 21 apps on 686 devices mostly happened on the first 6 days (Figure 5c) and the target apps were uninstalled within 3 days (Figure 5d), which is a very unusual synchronized behavior among a group of devices and apps. Moreover, in real Wi-Fi data, we find a group of students with a similar schedule, showing periodic and reasonable activities on the campus (Figure 6).

Reproducibility: Our code and datasets are publicly available at <https://github.com/BGT-M/AugSplicing>.

2 Related Work

Multi-aspect tuples can also be represented as attributed edges in a rich graph, e.g. users and objects as graph nodes, and rating scores and times as different attributes on graph edges. We, therefore, summarise the related research on dense block detection using both graphs and tensors (including two-mode matrices).

2.1 Static Tensors and Graphs

Dense subgraph detection has been extensively studied in (Hooi et al. 2016; Gibson, Kumar, and Tomkins 2005; Charikar 2000). Spectral decomposition based methods, e.g., SPOKEN (Prakash et al. 2010) considers the EigenSpokes on EE-plot produced by pairs of eigenvectors to detect near-cliques in social network. FRAUDAR (Hooi et al. 2016) considers both node and edge suspiciousness as a metric to detect frauds (i.e. dense blocks) and is also resistant to camouflage. CROSSPOT (Jiang et al. 2015) proposes an intuitive, principled metric satisfying the axioms that any metric of suspiciousness should obey, and design an algorithm to spot dense blocks sorted by order of importance (“suspiciousness”). HOSVD, CP Decomposition (CPD) (Kolda and Bader 2009) and disk-based algorithm (Oh et al. 2017) spot dense subtensors by Tensor decomposition. M-ZOOM (Shin, Hooi, and Faloutsos 2016) and D-CUBE (Shin et al. 2017a) adopt greedy approximation algorithms to detect dense subtensors with quality guarantees. CatchCore (Feng, Liu, and Cheng

2019) designs a unified metric optimized with gradient-based methods to find hierarchical dense subtensors. (Liu, Hooi, and Faloutsos 2018) optimizes the metric of suspiciousness from topology, rating time, and scores. ISG+D-spot (Yikun et al. 2019) constructs information sharing graph and finds dense subgraphs for the hidden-densest block patterns. Flock (Shah 2017) detects lockstep viewers in a live streaming platform.

However, these methods do not consider any temporal information, or only treat time bins as a static mode.

2.2 Dynamic Tensors and Graphs

In terms of dynamic graphs, some methods monitor the evolution of the entire graph and detect changes (density or structure) of subgraphs. SPOTLIGHT (Eswaran et al. 2018) utilizes graph sketches to detect the sudden density changes of a graph snapshot in a time period. SDRegion (Wong et al. 2018) detects blocks consistently becoming dense or sparse in genetic networks. EigenPulse (Zhang et al. 2019) is based on a fast spectral decomposition approach, single-pass PCA (Yu et al. 2017), to detect the density surges. Other methods, like MIDAS (Bhatia et al. 2020a,c) and MSTREAM (Bhatia et al. 2020b) detect suddenly arriving groups of suspiciously similar edges in edge streams, but do not take into account the topology of the graph. DENSESTREAM (Shin et al. 2017b) maintains dense blocks incrementally for every coming tuple and updates dense subtensors when it meets an updating condition, limiting the detection speed.

As for clustering-based methods, (Manzoor, Milajerdi, and Akoglu 2016) compare graphs based on the relative frequency of local substructures to spot anomalies. (Cao et al. 2014) uncovers malicious accounts that act similarly in a sustained period of time. Tensor decomposition-based methods, e.g., SamBaTen (Gujral, Pasricha, and Papalexakis 2018) and OnlineCP (Zhou et al. 2016) conduct the incremental tensor decomposition. Summarization based methods, e.g. (Shah et al. 2015) finds temporal patterns by summarizing important temporal structures. (Araujo et al. 2014) uses iterated rank-1 tensor decomposition, coupled with MDL (Minimum Description Length) to discover temporal communities.

Our method formulates the time-stamped tuples as a streaming tensor whose *time* mode is constantly augmented, such that the numerical value of entries in the previously observed tensor will not be changed. We incrementally splice incoming dense subtensors with the previous ones at each time step, achieving efficient results.

3 Definitions and Problem

We now give the notations used in the paper and describe our problem. Table 1 lists the key symbols.

Tensors are multi-dimensional arrays as the high-order generalization of vectors (1-dimensional tensors) and matrices (2-dimensional tensors). The number of dimensions of a tensor is its order, denoted by N . And each dimension is called a mode. For an N -mode tensor \mathcal{X} with non-negative entries, each (i_1, \dots, i_N) -th entry is denoted by $e_{i_1 \dots i_N}$. We use mode- n to indicate the n -th mode as (De Lathauwer, De Moor, and Vandewalle 2000; Lu, Plataniotis, and Venetsanopoulos 2013) do. Let i_n be mode- n index





Figure 1: (a) A tensor contains dense blocks (subtensors), and an incoming tensor at a time step contains tuples in range $(t, t + s]$. (b) AUGSPICING is fastest while maintaining similar accuracy (in terms of F-measure) compared to state-of-the-art approaches

Symbol	Definition
\mathcal{X}, \mathcal{B}	a tensor, and subtensor, i.e. block
$\mathcal{X}(t)$	N -mode tensor up to time t
N	number of modes in \mathcal{X}
$e_{i_1, \dots, i_N}(t)$	entry of $\mathcal{X}(t)$ with index i_1, \dots, i_N
$I_n(\cdot)$	set of mode- n indices of tensor
$M(\cdot)$	mass of tensor i.e. sum of non-zero entries
$S(\cdot)$	size of tensor
$g(\cdot)$	arithmetic degree density of tensor
s	augmenting time stride
$\mathcal{X}(t, s)$	N -mode augmenting tensor within time range $(t, t + s]$
k	number of blocks kept during iterations
$[x]$	$\{1, 2, \dots, x\}$

Table 1: Table of Symbols.

of entry $e_{i_1, \dots, i_n, \dots, i_N}$. We define the mass of \mathcal{X} as $M(\mathcal{X})$ to be the sum of its non-zero entries, and the size of \mathcal{X} as $S(\mathcal{X}) = \sum_{i=1}^N |I_i(\mathcal{X})|$, where $I_i(\mathcal{X})$ is the set of mode- i indices of \mathcal{X} . Let **block** \mathcal{B} be a subtensor of \mathcal{X} . Similarly, $M(\mathcal{B})$ and $S(\mathcal{B})$ are mass and size of block \mathcal{B} .

Our problem can be described as follows:

Problem 1 (Synchronized Behavior Detection in Streaming Tensor). *Given a stream of time-stamped tuples, i.e. streaming tensor $\mathcal{X}(t)$, and an augmenting time stride s , find top k dense blocks (i.e. subtensors) of $\mathcal{X}(t)$ so far at every tracking time step.*

We use the *arithmetic average mass* as the density measure of a block \mathcal{B} to avoid trivial solution as (Shin, Hooi, and Faloutsos 2016; Shin et al. 2017a), i.e. $g(\mathcal{B}) = \frac{M(\mathcal{B})}{S(\mathcal{B})}$.

4 Proposed Algorithm

In this section, we first theoretically analyze the splicing condition to increase density, and then guided by theory design a near-greedy algorithm to splice any two dense blocks. The overall algorithm (AUGSPICING) and time complexity are given in Section 4.3.

4.1 Theoretical Analysis for Splicing

We analyze the theoretical condition that whether splicing (i.e. merging partially) two dense blocks can result in a block with higher density, as Figure 2 shows. We call such merging as **splicing**.

Theorem 1 (Splicing Condition). *Given two blocks $\mathcal{B}_1, \mathcal{B}_2$ with $g(\mathcal{B}_1) \geq g(\mathcal{B}_2)$, $\exists \mathcal{E} \subseteq \mathcal{B}_2$ such that $g(\mathcal{B}_1 \cup \mathcal{E}) > g(\mathcal{B}_1)$ if and only if*

$$M(\mathcal{E}) > \sum_{n=1}^N r_n \cdot g(\mathcal{B}_1) = Q \cdot g(\mathcal{B}_1), \quad (1)$$

where $r_n = |I_n(\mathcal{E}) \setminus I_n(\mathcal{B}_1)|$, i.e. the number of new mode- n indices that \mathcal{E} brings into \mathcal{B}_1 . $Q = \sum_{n=1}^N r_n$, i.e. the total number of new indices that \mathcal{E} brings into \mathcal{B}_1 .

Proof. First, we prove the “ \Leftarrow ” condition. Based on the definition of $g(\cdot)$, we have

$$\begin{aligned} g(\mathcal{B}_1 \cup \mathcal{E}) &= \frac{M(\mathcal{B}_1) + M(\mathcal{E})}{S(\mathcal{B}_1) + Q} > \frac{M(\mathcal{B}_1) + Q \cdot g(\mathcal{B}_1)}{S(\mathcal{B}_1) + Q} \\ &= \frac{S(\mathcal{B}_1) \cdot g(\mathcal{B}_1) + Q \cdot g(\mathcal{B}_1)}{S(\mathcal{B}_1) + Q} = g(\mathcal{B}_1) \end{aligned}$$

Similarly, we can prove the “ \Rightarrow ” condition. ■

We can see that while splicing blocks, new indices of size Q are brought into some modes of \mathcal{B}_1 , and only merging the block \mathcal{E} with a large enough mass satisfying inequation (1), can increase $g(\mathcal{B}_1)$. Based on the theory, we design an effective algorithm to splice blocks as shown later.

4.2 Splicing Two Blocks

The purpose of splicing two blocks is to find a higher-density block by moving entries from one to another. Thus, based on the above analysis, the smaller size of new indices (i.e. smaller Q), and a larger mass of the merging block, can greedily increase the density of the spliced block. Our algorithm for splicing two given dense blocks is designed as Algorithm 1.

Algorithm 1 Splice two dense blocks

Input: two dense blocks: \mathcal{B}_1 and \mathcal{B}_2 , with $g(\mathcal{B}_1) \geq g(\mathcal{B}_2)$.

Output: new dense blocks

```

1: repeat
  /* minimize the size of new indices,  $Q$  */
2:    $\mathbf{q} \leftarrow$  get set of modes that have to bring new indices into  $\mathcal{B}_1$  for splicing
3:    $Q \leftarrow |\mathbf{q}|$   $\triangleright$  to minimize  $Q$ , considering only one new index from each mode in  $\mathbf{q}$ 
4:    $\mathbf{H} \leftarrow$  an empty max heap for blocks and ordered by block mass
5:   for each combination of new indices  $(i_{q_1}, \dots, i_{q_Q}), q \in \mathbf{q}$  do
6:      $\mathcal{E} \leftarrow$  block with entries  $\{e_{i_1 \dots i_{q_1} \dots i_{q_Q} \dots i_N} \in \mathcal{B}_2 \mid \forall n \in [N] \setminus \mathbf{q}, i_n \in I_n(\mathcal{B}_1)\}$ 
7:     push  $\mathcal{E}$  into  $\mathbf{H}$ 
8:   end for
  /* maximize  $M(\mathcal{E})$ , given  $Q$  */
9:   for  $\mathcal{E} \leftarrow \mathbf{H}.\text{top}()$  do
10:    if  $M(\mathcal{E}) > Q \cdot g(\mathcal{B}_1)$  then  $\triangleright$  inequation (1)
11:       $\mathcal{B}_1, \mathcal{B}_2 \leftarrow$  update  $\mathcal{B}_1 \cup \mathcal{E}, \mathcal{B}_2 \setminus \mathcal{E}$ 
12:      remove  $\mathcal{E}$  from  $\mathbf{H}$ , and re-heapify  $\mathbf{H}$ 
13:    else
14:      break
15:    end if
16:  end for
17: until no updates on  $\mathcal{B}_1$  and  $\mathcal{B}_2$ 
18: return new block  $\mathcal{B}'_1$  of higher density, and residual dense block  $\mathcal{B}_2$ 

```

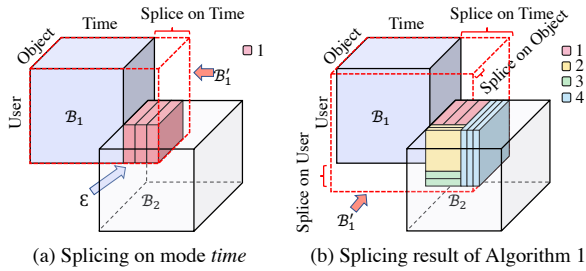


Figure 2: An illustration of splicing two blocks of three modes, i.e. (*user, object, time*). Given \mathcal{B}_1 and \mathcal{B}_2 with $g(\mathcal{B}_1) \geq g(\mathcal{B}_2)$ since there are no common indices on mode *time*, set $\mathbf{q} = \{\text{time}\}$ and $Q = 1$. After splicing red blocks 1 into \mathcal{B}_1 , all modes of two blocks are overlapped, and Algorithm 1 chooses one mode to bring new indices. As shown in (b), colored blocks 2, 3, 4 are successively spliced into \mathcal{B}_1 , bringing new indices into all three modes of \mathcal{B}_1 . \mathcal{B}'_1 is new \mathcal{B}_1 with higher density.

Given $g(\mathcal{B}_1) \geq g(\mathcal{B}_2)$, the idea is to alternatively find the minimum size of new indices (i.e. minimum Q), and maximum mass of block \mathcal{E} , given new indices for splicing.

In terms of Q , we first need to decide the set of modes \mathbf{q} that have to bring new indices into \mathcal{B}_1 (line 2). Since there are no common indices in mode- q of block \mathcal{B}_1 and \mathcal{B}_2 , at least one new index has to be added to the mode- q indices of \mathcal{B}_1 , then we add q to \mathbf{q} . Thus the minimum $Q = |\mathbf{q}|$. If all modes have common indices, then \mathbf{q} is empty and we do the following:

1. Let the block $\mathcal{E} \subseteq \mathcal{B}_2$ consist of entries of common indices.
2. Move non-zero entries of \mathcal{E} into \mathcal{B}_1 (if they exist), which

increases the density of \mathcal{B}_1 without bringing new indices.

3. Choose one mode q to splice. For each mode of $[N]$, we generate subblocks of \mathcal{B}_2 by choosing one new index on this mode, and all indices overlapped with \mathcal{B}_1 on other modes. Subblock with maximum mass was generated from mode q . In such a scenario, $Q = 1$ to choose only one mode to splice.

For mass maximization, we use a max heap to organize blocks by their mass (line 4). Top of the max heap is always the maximum mass block. Then we enumerate all possible combinations of a new index from each mode in \mathbf{q} (lines 5-7) to build a max heap \mathbf{H} . Since the number of modes of blocks, N , is usually as small as $3 \sim 5$ for real data, and the size of possible combinations is comparable to $S(\mathcal{B}_2)$, given \mathcal{B}_2 is a small-size block in original tensor \mathcal{X} . Moreover, according to inequation (1), only those blocks with large enough masses are add into max heap \mathbf{H} . Then we splice a maximum-mass block on top of \mathbf{H} , iteratively increasing $g(\mathcal{B}_1)$ and merging next top block satisfying $M(\mathcal{E}) > Q \cdot g(\mathcal{B}_1)$, until no large-mass blocks remain for merging (lines 8-13).

With first getting the minimum size of new indices, i.e. minimum Q , and constantly merging maximum-mass block by choosing new indices into \mathcal{B}_1 , our algorithm ends until no updates can be made on \mathcal{B}_1 and \mathcal{B}_2 .

Example 1. Figure 2 gives a running example of our algorithm. In the beginning, \mathcal{B}_1 and \mathcal{B}_2 have no common indices on mode *time*, thus $\mathbf{q} = \{\text{time}\}$. Alg 1 splices on mode *time* with red blocks 1 merged into \mathcal{B}_1 , forming new \mathcal{B}'_1 of higher density (i.e. \mathcal{B}'_1 in Figure 2(a)). Note that each red block brings only one new index into \mathcal{B}_1 , i.e. $Q = 1$. Afterward, all modes of two new blocks have common indices. Since \mathcal{B}_2 doesn't have any non-zero entry of common indices with

\mathcal{B}_1 , Alg 1 has to choose one mode q to bring new indices into \mathcal{B}_1 for splicing. q is successively the mode object, time, user in the example. In the end, colored blocks **1, 2, 3, 4** are successively merged. A new block \mathcal{B}_1 (i.e. \mathcal{B}'_1) with higher density, and residual block \mathcal{B}_2 are returned (Figure 2(b)).

4.3 Overall Algorithm

In this section, we first describe the overall algorithm for incrementally detecting the densest blocks at each time step $t + s$, then analyze the time complexity of AUGSPlicing, which is near-linear with the number of non-zero tuples.

Let bold symbols $\mathcal{B}(t)$ and $\mathcal{C}(t)$ be sets of top $k + l$ dense blocks of previous $\mathcal{X}(t)$ and incoming $\mathcal{X}(t, s)$, where l is a slack constant for approximating the top k dense blocks with l more blocks. Our overall algorithm is as follows:

(a) **Splice two dense blocks:** We iteratively choose two candidate blocks from $\mathcal{B}(t) \cup \mathcal{C}(t)$, denoted as \mathcal{B}_1 and \mathcal{B}_2 with $g(\mathcal{B}_1) \geq g(\mathcal{B}_2)$, then use Algorithm 1 to splice them.

(b) **Iteration and Output:** The splicing iteration stops until no blocks can be spliced, or reach the maximum number of epochs. Then, AUGSPlicing outputs top k of $k + l$ dense blocks at time step $t + s$, and moves on to the next time step with $k + l$ blocks.

Theorem 2 (Time Complexity). *The time complexity of AUGSPlicing at time step $t + s$ is $O(N^2(k + l)nnz(\mathcal{X}(t, s))L(\mathcal{X}(t, s)) + 2(k + l)nnz(\mathcal{B}) \log S(\mathcal{B}))$, where $L(\cdot) = \max_{n \in [N]} |I_n(\cdot)|$ and $nnz(\cdot)$ is the number of non-zero entries.*

Proof. At time step $t + s$, the incoming tensor is $\mathcal{X}(t, s)$, and the complexity for detecting new top $k + l$ dense blocks is $O(N^2(k + l)nnz(\mathcal{X}(t, s))L(\mathcal{X}(t, s)))$ according to [20].

Let \mathcal{B} be a block of the maximum non-zero entries, and the largest size among splicing blocks. Compared to building a max heap, a more time-consuming procedure is the iteration of updating and re-heapifying when new entries are merged into a block. Considering the worst case that all the blocks are merged into one, at most $2(k + l)nnz(\mathcal{B})$ entries are spliced, i.e. the maximum number of updates in the max heap. Therefore the time complexity for iterative splices is at most $O(2(k + l)nnz(\mathcal{B}) \log S(\mathcal{B}))$, as the heap size is $O(S(\mathcal{B}))$. Thus the complexity of AUGSPlicing at time step $t + s$ is $O(N^2(k + l)nnz(\mathcal{X}(t, s))L(\mathcal{X}(t, s)) + 2(k + l)nnz(\mathcal{B}) \log S(\mathcal{B}))$. ■

Since $nnz(\mathcal{B}) = O(nnz(\mathcal{X}(t, s)))$ for proper stride s , our algorithm is near-linear in the number of incremental tuples $nnz(\mathcal{X}(t, s))$ in practice as (Shin et al. 2017a) shows, which ensures near-linear in the number non-zero entries of streaming tensors.

5 Experiments

We design the experiments to answer the following questions:

Q1. Speed and Accuracy: How fast and accurate does our algorithm run compared to the state-of-the-art streaming algorithms and the re-run of batch algorithms on real data?

Q2. Real-World Effectiveness: Which anomalies or lock-step behavior does AUGSPlicing spot in real data?

Name	Volume	# Edges
Rating data (user, item, timestamp)		
Yelp	468K × 73.3K × 0.73K	1.34M
BeerAdvocate	26.5K × 50.8K × 0.5K	1.08M
mobile devices, app, installing time, uninstalling time		
App	2.47M × 17.9K × 30 × 30	5.43M
device IP, Wi-Fi AP, connecting time, disconnecting time		
Wi-Fi	119.4K × 0.84K × 1.46K × 1.46K	6.42M

Table 2: Data Statistics.

Q3. Scalability: How does the running time of our algorithm increase as input tensor grows?

Datasets: Table 2 lists the real-world data used in our paper. Two rating data are publicly available. App data is mobile device-app installation and uninstallation data under an NDA agreement from a company. Wi-Fi data is device-AP connection and disconnection data from Tsinghua University.

Experimental Setup: All experiments are carried out on a 2.3GHz Intel Core i5 CPU with 8GB memory. We compare our method with state-of-the-art streaming dense block detection method, DENSESTREAM, and the re-run of batch methods, D-CUBE, CROSSPOT, and CP Decomposition (CPD). D-CUBE is implemented in Java to detect dense blocks in tensor $\mathcal{X}(t, s)$. Specifically, we use “arithmetic average mass” as the metric of D-CUBE. We use a variant of CROSSPOT which maximizes the same metric and use the CPD result for seed selection similar to (Shin, Hooi, and Faloutsos 2016). We set time stride s to 30 in a day for Yelp, 15 in a day for BeerAdvocate, 1 in a day for App and Wi-Fi, as different time granularity. k is set to 10 and l to 5 for all datasets.

5.1 Speed and Accuracy

For detecting dense blocks in a streaming setting, our method only deals with augmenting tensor with *stride* size at each time step, then it combines detected blocks of the incremental tensor with the previous results to detect dense blocks until now. In contrast, the batch algorithms are re-run for the holistic tensor from scratch at each time step to detect dense blocks of augmenting tensors. DENSESTREAM needs to maintain a dense subtensor when a new entry comes, which is very time-consuming. We measure the wall-clock time taken by each method and the results are as shown in Figure 3a. As we can see, AUGSPlicing is the fastest. It is 320× faster than DENSESTREAM, 1.8× faster than D-CUBE, 3.2× faster than CROSSPOT and 13× faster than CPD on Wi-Fi dataset.

To demonstrate the accuracy of our algorithm, we track the density of the densest block found by AUGSPlicing and other methods while the tensor augments at each time step as (Shin et al. 2017b) does and the result is shown in Figure 4a-4d. We can see that the densest block has close

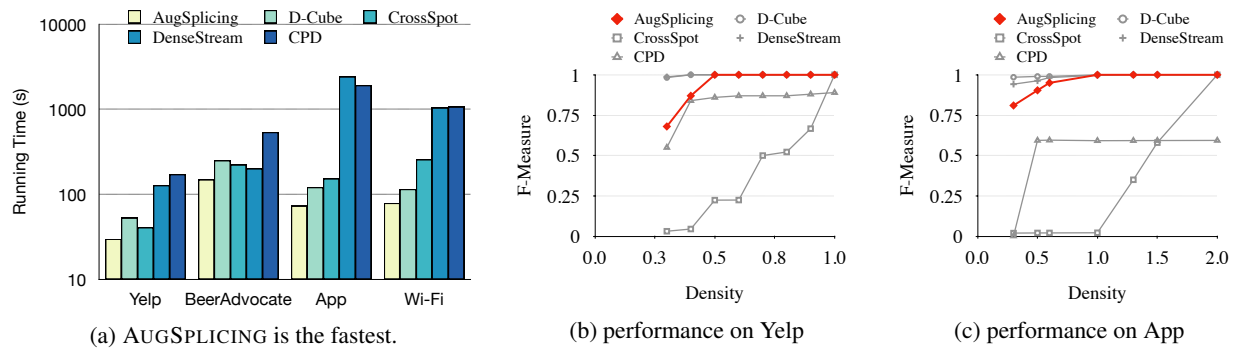


Figure 3: AUGSP LICING is fast and accurate. (a) AUGSP LICING is 320 \times faster than baselines. In (b)-(c), our method has accuracy (F-measure) comparable to the state-of-the-art methods: DENSESTREAM and D-CUBE, especially when injected fraudulent density is larger than 0.5.




Figure 4: AUGSP LICING has comparable accuracy (density) with the state-of-the-art methods.

density to that found by DENSESTREAM and the re-run of D-CUBE for long time steps, though accumulated error.

We now explain why AUGSP LICING achieves comparable high accuracy. Due to the skewness of real graphs, densities of top dense blocks can be very skewed, which reduces the probability of the top k dense blocks of $\mathcal{X}(t+s)$ having overlapped nodes with top $(k+l)$ or lower dense blocks in $\mathcal{X}(t)$ and $\mathcal{X}(t,s)$. Due to the principle of time locality, tuples of dense blocks will be close in mode *time*. Thus AUGSP LICING can detect top k dense blocks with comparable high density by sufficient splices.

Detection in injected attacks: For Yelp data, we injected 100 fraudulent users and items in a week with the volume density ranging from 1 to 0.1. For app data, an app’s rank is decided by its downloads, which improves by 1 if the app is installed and remains more than the required amount of days by a mobile device in a real scenario. Then we injected 500 apps and 5000 mobile devices, with the installing time uniformly distributed in 3 days. The uninstalling time was the installing time plus a week with the volume density ranging from 2.0 to 0.1. Intuitively, the smaller the density of injected blocks, the harder it is to detect, and the block with a density of 0.1 is quite difficult to detect. Figures 3b-3c show that F-measure of AUGSP LICING increases rapidly as density increases from 0.3 to 0.5 and remains higher than 90% when the density reaches to 0.5, presenting our ability in detecting fraudulent mobile devices and apps.

5.2 Effectiveness

Results on App data with ground-truth labels: In this section, we verify that AUGSP LICING accurately detects a dense block of fraudulent accounts in App data, as verified by clear signs of fraud exhibited by a majority of detected mobile devices and apps. We collected the devices detected by all methods and manually labeled by the company who owns App data, based on empirical knowledge on features of devices: e.g. locations, the number of installed apps, and the number of apps opened for the first time in a day, etc. For example, devices are identified as fraudulent if they appear in dozen cities in a day.

Figure 1b shows both accuracy (F-measure) and speed (elapsed running time) for all comparison methods. We can see that the AUGSP LICING runs 320 \times faster than the state-of-the-art streaming algorithm, DENSESTREAM, keeping comparable accuracy. Compared to the fast re-run D-CUBE, AUGSP LICING achieves 1.8 \times faster, and much higher accuracy. Figures 5a-5d present the detailed information of the densest block detected by AUGSP LICING. We draw “degree” distributions for detected 686 devices and 21 apps in Figure 5a and 5b. Note that the “degree” of a mobile device is # of apps installed by the mobile. Similarly “degree” of an app is the number of devices installing the app. As a result, 365 devices from 686 detected devices have been identified as fraudsters by the company, which is a very high concentration in a fraud detection task, considering a small fraction of




Figure 5: (a) and (b) show AUGSP LICING detects a real-world suspicious block that is explainable: 686 mobile devices repeatedly installed and uninstalled 21 apps 5.66×10^4 times in total, which is very unusual for a group of devices and apps. (c) and (d) show an 8-day installing time period and all that suspicious apps were uninstalled within one week, and most of them stayed only up to 3 days on a suspicious device.




Figure 6: AUGSP LICING finds a big community of students having synchronized busy schedule in Wi-Fi data. Wi-Fi connections reached the peak at around 10 AM every day (red dash lines) in (c). Most of the connections stayed around 1 to 2 hours as shown in (d).

fraudulent devices over the whole devices. Actually, devices not identified as fraudsters are very suspicious by analyzing their behavior: 686 mobile devices repeatedly installed and uninstalled 21 apps 5.66×10^4 times, with up to 7100 times for one app. Furthermore, all the installations were concentrated in a short time period (i.e. 8 days) and uninstalled within one week afterward (see Figures 5c-5d). It is very likely that these mobile devices boost these apps’ ranks in an app store by installations and uninstallations in lockstep.

Results on Wi-Fi data: We discover synchronized patterns that may interest student administrators. Figure 6 shows the densest block detected by AUGSP LICING in Wi-Fi data. Figure 6a and 6b show 2849 devices and 413 Wi-Fi APs which had 8.03×10^5 connections/disconnections in total, indicating a busy schedule for many students. As shown in Figure 6c, behavior of this group of students was periodic and synchronized. Wi-Fi connections concentrated from 8 AM - 5 PM every day and reached a peak at around 10 AM (red dotted line). That may be because students’ first and second classes begin around 8 AM and 10 AM respectively. Moreover, Figure 6d shows that most of the connections stayed around 1-2 hours, which is the usual length of one class.

5.3 Scalability

Figure 7a and 7b show that the running time of our algorithm scales linearly with the size of non-zero entries up to the current time step, consistent with the time complexity analysis result in Section 4.3, while the running time of the re-run of D-CUBE is quadratic in streaming tensors.

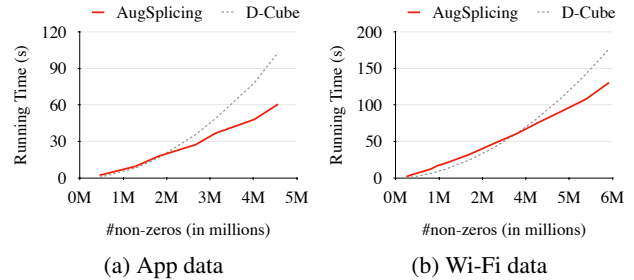


Figure 7: AUGSP LICING has near-linear running time.

6 Conclusions

In this paper, we model a stream of tuples as a streaming tensor, and propose a streaming algorithm, AUGSP LICING, to spot the most synchronized behavior which indicates anomalies or interesting patterns efficiently and effectively. Our main contributions are as follows:

1. **Fast and streaming algorithm:** Our approach can effectively capture synchronized activity in streams, up to $320 \times$ faster than the best streaming method (Figure 1b).
2. **Robustness:** Our method is robust with theory-guided incremental splices for dense block detection.
3. **Effectiveness:** Our method is able to detect anomalies and discover interesting patterns accurately in real-world data (Figures 5a-5d and Figures 6a-6d).

Ethical Impact

We contribute a fast incremental algorithm to detect dense blocks formed by synchronized behavior in a stream of time-stamped tuples. Our work has wide applications on anomaly detection tasks, e.g. app store fraud detection (e.g. suspicious mobile devices boosting target apps' ranks in recommendation list in an app store), rating fraud detection in review sites and etc. In addition, our work can be applied to discover interesting patterns or communities in real data (e.g. revealing a group of students having the same classes of interest). Our approach can scale to very large data, update the estimations when data changes over time efficiently, and incorporate all the information effectively. Our work is even more applicable to online data mining tasks, especially when large-sized data arrives at a high rate. While most experiments are cybersecurity related, one experiment detects student communities from Wi-Fi data. From a societal impact, potential misuse against privacy has to be taken care of.

Acknowledgments

This paper is partially supported by the National Science Foundation of China under Grant No.U1911401, 61772498, 61872206, 91746301. This paper is also supported by the Strategic Priority Research Program of the Chinese Academy of Sciences, Grant No. XDA19020400, National Key R&D Program of China (2020AAA0103502) and 2020 Tencent Wechat Rhino-Bird Focused Research Program.

References

- Akoglu, L.; Tong, H.; and Koutra, D. 2015. Graph based anomaly detection and description: a survey. *Data mining and knowledge discovery* 29(3): 626–688.
- Araujo, M.; Papadimitriou, S.; Günnemann, S.; Faloutsos, C.; Basu, P.; Swami, A.; Papalexakis, E. E.; and Koutra, D. 2014. Com2: fast automatic discovery of temporal ('comet') communities. In *PAKDD*, 271–283. Springer.
- Bhatia, S.; Hooi, B.; Yoon, M.; Shin, K.; and Faloutsos, C. 2020a. Midas: Microcluster-Based Detector of Anomalies in Edge Streams. In *AAAI*, 3242–3249.
- Bhatia, S.; Jain, A.; Li, P.; Kumar, R.; and Hooi, B. 2020b. MStream: Fast Anomaly Detection in Multi-Aspect Streams. *arXiv preprint arXiv:2009.08451*.
- Bhatia, S.; Liu, R.; Hooi, B.; Yoon, M.; Shin, K.; and Faloutsos, C. 2020c. Real-Time Streaming Anomaly Detection in Dynamic Graphs. *arXiv preprint arXiv:2009.08452*.
- Cao, Q.; Yang, X.; Yu, J.; and Palow, C. 2014. Uncovering large groups of active malicious accounts in online social networks. In *Proc. of the 2014 ACM SIGSAC*, 477–488. ACM.
- Charikar, M. 2000. Greedy approximation algorithms for finding dense components in a graph. In *International Workshop on Approximation Algorithms for Combinatorial Optimization*, 84–95. Springer.
- De Lathauwer, L.; De Moor, B.; and Vandewalle, J. 2000. On the best rank-1 and rank-(r_1, r_2, \dots, r_n) approximation of higher-order tensors. *SIAM journal on Matrix Analysis and Applications* 21(4): 1324–1342.
- Eswaran, D.; Faloutsos, C.; Guha, S.; and Mishra, N. 2018. SpotLight: Detecting Anomalies in Streaming Graphs. In *SIGKDD*, 1378–1386. ACM.
- Feng, W.; Liu, S.; and Cheng, X. 2019. CatchCore: Catching Hierarchical Dense Subtensor. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 156–172. Springer.
- Gibson, D.; Kumar, R.; and Tomkins, A. 2005. Discovering large dense subgraphs in massive graphs. In *Proc. of the 31st VLDB*, 721–732. VLDB Endowment.
- Gujral, E.; Pasricha, R.; and Papalexakis, E. E. 2018. Sambaten: Sampling-based batch incremental tensor decomposition. In *SDM*, 387–395. SIAM.
- Hooi, B.; Song, H. A.; Beutel, A.; Shah, N.; Shin, K.; and Faloutsos, C. 2016. Fraudar: Bounding graph fraud in the face of camouflage. In *Proc. of the 22nd ACM SIGKDD*, 895–904. ACM.
- Jiang, M.; Beutel, A.; Cui, P.; Hooi, B.; Yang, S.; and Faloutsos, C. 2015. A general suspiciousness metric for dense blocks in multimodal data. In *2015 IEEE International Conference on Data Mining*, 781–786. IEEE.
- Jiang, M.; Cui, P.; Beutel, A.; Faloutsos, C.; and Yang, S. 2014. Catchsync: catching synchronized behavior in large directed graphs. In *Proc. of the 20th ACM SIGKDD*, 941–950. ACM.
- Kolda, T. G.; and Bader, B. W. 2009. Tensor decompositions and applications. *SIAM review* 51(3): 455–500.
- Liu, S.; Hooi, B.; and Faloutsos, C. 2018. A contrast metric for fraud detection in rich graphs. *IEEE Transactions on Knowledge and Data Engineering* 31(12): 2235–2248.
- Lu, H.; Plataniotis, K. N.; and Venetsanopoulos, A. 2013. *Multilinear subspace learning: dimensionality reduction of multidimensional data*. CRC press.
- Manzoor, E.; Milajerdi, S. M.; and Akoglu, L. 2016. Fast memory-efficient anomaly detection in streaming heterogeneous graphs. In *Proc. of the 22nd ACM SIGKDD*, 1035–1044. ACM.
- Oh, J.; Shin, K.; Papalexakis, E. E.; Faloutsos, C.; and Yu, H. 2017. S-hot: scalable high-order tucker decomposition. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, 761–770. ACM.
- Prakash, B. A.; Sridharan, A.; Seshadri, M.; Machiraju, S.; and Faloutsos, C. 2010. Eigenspokes: Surprising patterns and scalable community chipping in large graphs. In *PAKDD*. Springer.
- Shah, N. 2017. FLOCK: Combating astroturfing on livestreaming platforms. In *Proc. of the 26th International Conference on World Wide Web*, 1083–1091. International World Wide Web Conferences Steering Committee.
- Shah, N.; Koutra, D.; Zou, T.; Gallagher, B.; and Faloutsos, C. 2015. Timecrunch: Interpretable dynamic graph summarization. In *Proc. of the 21th ACM SIGKDD*, 1055–1064. ACM.

- Shin, K.; Hooi, B.; and Faloutsos, C. 2016. M-zoom: Fast dense-block detection in tensors with quality guarantees. In *ECML/PKDD*, 264–280. Springer.
- Shin, K.; Hooi, B.; Kim, J.; and Faloutsos, C. 2017a. D-cube: Dense-block detection in terabyte-scale tensors. In *Proc. of the Tenth WSDM*, 681–689. ACM.
- Shin, K.; Hooi, B.; Kim, J.; and Faloutsos, C. 2017b. Densealert: Incremental dense-subtensor detection in tensor streams. In *Proc. of the 23rd ACM SIGKDD*, 1057–1066. ACM.
- Teng, S.-H. 2016. Scalable algorithms for data and network analysis. *Foundations and Trends® in Theoretical Computer Science* 12(1–2): 1–274.
- Wong, S. W.; Pastrello, C.; Kotlyar, M.; Faloutsos, C.; and Jurisica, I. 2018. SDREGION: Fast Spotting of Changing Communities in Biological Networks. In *Proc. of the 24th ACM SIGKDD*, 867–875. ACM.
- Yikun, B.; Xin, L.; Ling, H.; Yitao, D.; Xue, L.; and Wei, X. 2019. No Place to Hide: Catching Fraudulent Entities in Tensors. In *The World Wide Web Conference*, 83–93. ACM.
- Yu, W.; Gu, Y.; Li, J.; Liu, S.; and Li, Y. 2017. Single-pass PCA of large high-dimensional data. In *IJCAI*, 3350–3356.
- Zhang, J.; Liu, S.; Yu, W.; Feng, W.; and Cheng, X. 2019. EigenPulse: Detecting Surges in Large Streaming Graphs with Row Augmentation. In *PAKDD*, 501–513. Springer.
- Zhou, S.; Vinh, N. X.; Bailey, J.; Jia, Y.; and Davidson, I. 2016. Accelerating online cp decompositions for higher order tensors. In *Proc. of the 22nd ACM SIGKDD*, 1375–1384. ACM.