

# TENSORSPLAT: Spotting Latent Anomalies in Time

Danai Koutra, Evangelos E. Papalexakis, Christos Faloutsos  
School of Computer Science,  
Carnegie Mellon University,  
Pittsburgh, PA, USA

Email: {danai, epapalex, christos}@cs.cmu.edu

**Abstract**—How can we spot anomalies in large, time-evolving graphs? When we have multi-aspect data, e.g. who published which paper on which conference and on what year, how can we combine this information, in order to obtain good summaries thereof and unravel hidden anomalies and patterns? Such multi-aspect data, including time-evolving graphs, can be successfully modelled using Tensors.

In this paper, we show that when we have multiple dimensions in the dataset, then tensor analysis is a powerful and promising tool. Our method TENSORSPLAT, at the heart of which lies the “PARAFAC” decomposition method, can give good insights about the large networks that are of interest nowadays, and contributes to spotting micro-clusters, changes and, in general, anomalies. We report extensive experiments on a variety of datasets (co-authorship network, time-evolving DBLP network, computer network and Facebook wall posts) and show how tensors can be proved useful in detecting “strange” behaviors.

## I. INTRODUCTION

Which authors are highly inter-disciplinary and which changed research interests during their career? Was there an attempt for network intrusion? When did a port-scan happen? In the data mining community, the last years, there is a surge of interest in detecting anomalous behaviors like the above mentioned in various, large datasets.

The list of applications of anomaly or outlier detection is huge: spammers in blogs, telemarketers in phone-call datasets, exfiltrators in companies, fraud detection in bank accounts or credit cards [7], malware detection [13], voting irregularity analysis, network intrusions [25], data cleansing etc.

The interest lies not only in detecting “strange” or “anomalous” behaviors, but also rare events and changes. For instance, a big number of wall postings in Facebook may be an indicator of birthday or nameday, if the duration of the event is small (e.g., one day); on the other hand, if from that day on the number of wall postings is bigger than usual, then the event might suggest that the account was compromised. In both cases, it is worth having tools that can detect and explain the change (even if an analyst with domain knowledge is required in the loop).

In a slightly different pace, suppose that we discover an author who, up to a certain point in time, used to publish to conference  $A$ ; if we discover that after a certain date, this author has switched from publishing to conference  $A$  and instead is associated with conference  $B$ , this may indicate that this author has switched research areas. The list of anomaly detection application continues: consider network traffic measurements that record the source-destination pairs,

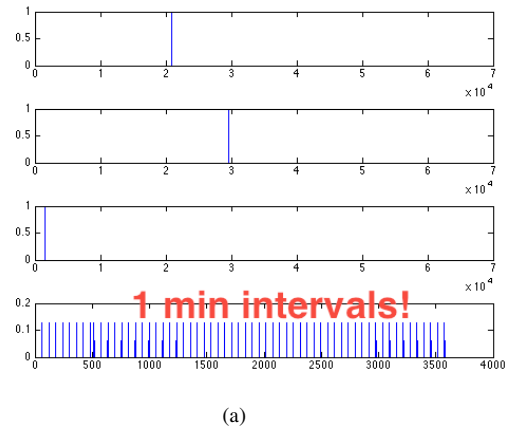


Fig. 1. Anomaly spotted on the LBNL dataset: bot attack-like behavior on port 1544, where the activity spikes in the temporal domain occur precisely every one minute.

the number of packets transmitted, the connection port, as well as the timestamp when the connection was initiated. It is thus of fundamental interest to network administrators to detect suspicious behaviors over this set of connections, such as a sequential port-scanning, which may well signify an impending attack and should definitely raise a flag to the administrator. A snapshot of our interesting findings may be found in Figure 1, which illustrates what appears to be a bot attack-like behavior on network traffic data.

The increasing interest in anomaly detection has led to the development of several methods that specialize either in multi-dimensional points or networks [22]. One of the most prevalent techniques is statistical analysis of the data [11]; the goal is to find patterns and laws or fit distributions to the data. The points that do not follow the distribution or do not obey the expected patterns are considered anomalous and require further investigation. Recently developed methods applied to networks rely on extracting features for the nodes of the graphs and converting the graph data to multi-dimensional points [2]. Then classic algorithms [9] [24], which are analyzed in the survey, are used to find distant or low-density-neighborhood points which may correspond to anomalous nodes.

In this work, we propose the detection of micro-clusters, rare events and changes in behaviors using tensors. Tensors are powerful tools that can be used for the analysis of the increasing in size data which contains multiple dimensions;

one of the most tangible examples where tensors can be used nicely is the case of time-evolving graphs. Our contributions are the following:

- We focus on how to use TENSORSPLAT, which is based on “PARAFAC”, a highly interpretable tensor decomposition method, in order to spot anomalies in data, and
- Report results in 3 different settings: (a) a small subgraph of the co-authorship DBLP network, (b) the time-evolving DBLP network that spans 49 years, and (c) LBNL, a big network traffic dataset.

The following sections are organized in the usual way: we first explain the theoretical concepts of our proposed method, present the experimental results and describe the related work.

## II. BACKGROUND

In this section, we provide the theoretical background required, in order to get a grasp of tensors and tensor decompositions. A very concise and comprehensive tutorial about tensors may be found in [19]. Table I presents the notation that we use in this section.

TABLE I  
DESCRIPTION OF MAJOR SYMBOLS.

Notation	Example	Description
italic, lowercase letters	$a$	scalar
bold, lowercase letters	$\mathbf{x}$	column vector
bold capital letters	$\mathbf{A}$	matrix
Euler script letters	$\mathcal{X}$	tensor
superscript	$\mathbf{x}^{(I)}$ , $\mathbf{X}^{(I \times J)}$ , $\mathcal{X}^{(I \times J \times K)}$	dimensions of vector, matrix, or tensor
index	$\mathbf{x}^{(i)}$ , $\mathbf{X}^{(i, j)}$ , $\mathcal{X}^{(i, j, k)}$	corresponding element of vector, matrix, or tensor

### A. Tensors

An  $n$  mode/way tensor is essentially a structure that is indexed by  $n$  variables. In particular, a 2-way tensor is nothing but a regular matrix, whereas a 3-way tensor may be viewed as a data cube. For the purposes of background exposition, we are going to focus on 3-way tensors, because they are easier to visualize; however, everything discussed in the sequel readily extends to higher ways, and in fact, one of our data mining case studies operates on a 4-way tensor.

### B. Matrix decomposition

In order to demonstrate the concept of a tensor decomposition, we will at first diverge to describe how one may decompose a matrix and how this relates to well established concepts in Information Retrieval and Data Mining. In particular, consider an  $I \times J$  matrix  $\mathbf{X}$ ; for instance, this matrix may correspond to an author by conference matrix, where the  $(i, j)$ -th entry indicates that author  $i$  has published a certain number of papers in conference  $j$ .

First, we need to define the Singular Value Decomposition (SVD) of  $\mathbf{X}$ , which states that any matrix  $\mathbf{X}$  may be decomposed in the following way

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$$

where  $\mathbf{U}$ ,  $\mathbf{V}$  are orthonormal matrices and  $\mathbf{\Sigma}$  is diagonal, with non-negative values on the diagonal, the so-called singular values of  $\mathbf{X}$ . In [16], it is shown that if we truncate the SVD of  $\mathbf{X}$  to a rank  $f$  lower than the actual rank of the matrix, this yields the optimal low rank approximation of the matrix, in the least squares sense.

If we call matrix  $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}$  and matrix  $\mathbf{B} = \mathbf{V}$  then we may rewrite the decomposition as:

$$\mathbf{X} \approx \mathbf{a}_1 \mathbf{b}_1^T \cdots \mathbf{a}_F \mathbf{b}_F^T$$

where equality holds if  $F = \text{rank}(\mathbf{X})$ . The above expression is essentially a *bilinear decomposition* of  $\mathbf{X}$ . We chose to obtain this bilinear decomposition through the SVD, but in fact, there exist numerous different approaches, e.g. the Non-Negative Matrix Factorization [20].

At this point, one may wonder what is the practicality of the above formulation. The answer was initially given in [15], where Latent Semantic Indexing (LSI) is introduced. In a nutshell, what LSI suggests is the following: Consider the author-by-conference example that we mentioned. If we take a rank  $f$  approximation of this matrix, then, in a sense, we force all authors and all conferences to be expressed using a basis of  $f$  vectors. In this way, we are grouping together authors and conferences that are expressed similarly in the matrix, i.e. they are “close”, usually, with respect to euclidean distance.

### C. Tensor decomposition

Having introduced the SVD and the bilinear decomposition, we are ready to extend our paradigm to the tensor regime. However, there is no single decomposition that fully extends the matrix SVD; we choose to elaborate on the most intuitive and easily interpretable of all, the PARAFAC decomposition.

Consider a three way tensor  $\mathcal{X}$  of dimensions  $I \times J \times K$ ;

The PARAFAC [17], [10] (also known as CP or trilinear) decomposition of  $\mathcal{X}$  in  $F$  components is an immediate extension of the bilinear decomposition, for tensors, i.e.

$$\mathcal{X} \approx \sum_{f=1}^F \mathbf{a}_f \circ \mathbf{b}_f \circ \mathbf{c}_f$$

The three way outer product of vectors  $\mathbf{a}$ ,  $\mathbf{b}$ ,  $\mathbf{c}$  is defined as

$$[\mathbf{a} \circ \mathbf{b} \circ \mathbf{c}](i, j, k) = \mathbf{a}(i)\mathbf{b}(j)\mathbf{c}(k)$$

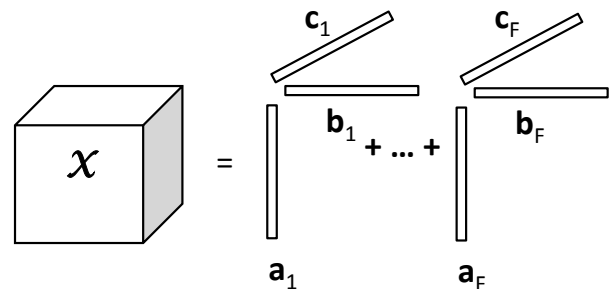


Fig. 2. The PARAFAC decomposition of  $\mathcal{X}$ .

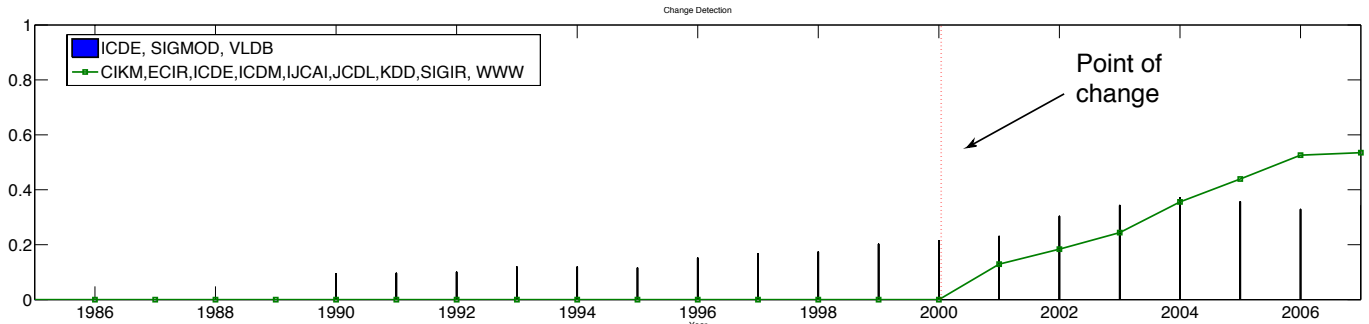


Fig. 3. In this Figure we demonstrate how TENSORSPLAT is able to perform change detection. In particular, we observe two components in which a well-known professor appears as an author; the first component mainly contains Databases conferences, whereas the second contains Data Mining conferences. The dashed red line indicates the point of change in research direction.

More compactly, we can write the PARAFAC decomposition as a triplet of matrices  $\mathbf{A}$ ,  $\mathbf{B}$ ,  $\mathbf{C}$ , i.e.

$$\mathbf{x} \approx [\mathbf{A}^{(I \times F)}, \mathbf{B}^{(J \times F)}, \mathbf{C}^{(K \times F)}]$$

the  $f$ -th column of which contains  $\mathbf{a}_f$ ,  $\mathbf{b}_f$  and  $\mathbf{c}_f$  respectively. One could normalize each column of the three factor matrices, and introduce a scalar term  $\lambda_f$ , one for each rank-one factor of the decomposition. Both representations, however, are equivalent.

#### D. Extending to higher ways

As we said before, we can have tensors with more than 3 ways/modes. For instance, if we had a 4-way tensor, then we would need to introduce a fourth vector  $\mathbf{d}$  on the outer product, and hence, the PARAFAC decomposition of tensor  $\mathbf{X}$  of dimensions  $I \times J \times K \times N$  would be

$$\mathbf{X} \approx \sum_{f=1}^F \mathbf{a}_f \circ \mathbf{b}_f \circ \mathbf{c}_f \circ \mathbf{d}_f$$

Furthermore, we would introduce the  $N \times F$  matrix  $\mathbf{D}$  whose columns contain the  $\mathbf{d}_f$  vectors.

#### E. Other Decompositions

Apart from the PARAFAC decomposition, there exist numerous other decompositions; a detailed exposition thereof may be found in [19].

### III. DATA ANALYSIS

TENSORSPLAT is our proposed method for anomaly or change detection in graphs where more than 2 modes are available (e.g., we have re-tweet activity between twitter users over time, or links between pages with context information). In these cases, a matrix does not suffice to completely describe the available data; a more sophisticated structure is needed. Our experiments show that tensors, a generalization of matrices to more dimensions, is a structure worth using when problems of this kind arise.

TENSORSPLAT is essentially based on the ‘‘PARAFAC’’ decomposition and a modified version of the Matlab Tensor

Toolbox [4], [6] which accommodates the needs of the applications we study. Specifically, we do ‘‘truncated’’ tensor decomposition by thresholding the small values that have little interest in the context of anomalies. Among the multiple components of the decomposition, we pick the most interesting ones, and, in this section, we elaborate on them.

The TENSORSPLAT method has three thrusts:

- 1) Change detection over time: We detect authors who have switched research areas, by observing changes in their publishing behavior.
- 2) Anomaly detection over time: We spot suspicious behaviors in a large set of computer network connections and in social network interactions.
- 3) Clustering: We identify clusters of authors with similar research interests and background who publish paper on relevant conferences.

The strength of TENSORSPLAT is attributed to its ability to accomplish the aforementioned tasks while operating on large amounts of data, which prove prohibitive for eye-balling, empirical analysis, an approach that would be literally looking for a needle in a haystack.

In the following lines, we provide a detailed description of both the data analyzed and the methods used.

#### A. Data

TABLE II  
LIST OF DATASETS

Name	Description	Dimensions
DBLP-1	(author, paper, conf)	$14.5K \times 14.4K \times 20$
DBLP-2	(author, conf, year)	$418K \times 3.5K \times 49$
LBNL	(src, dst, port #, time)	$65K \times 65K \times 65K \times 3.6K$
FACEBOOK	(wall, poster, day)	$64K \times 64K \times 1.8K$

Before we delve into the details of our experimental setup and results explanation, we will describe the datasets we used (Table II).

- DBLP-1 [8]: The small DBLP sub-network consists of 14,376 papers, 14,475 authors and 20 conferences in the areas of Data Mining (DM), Databases (DB), Information

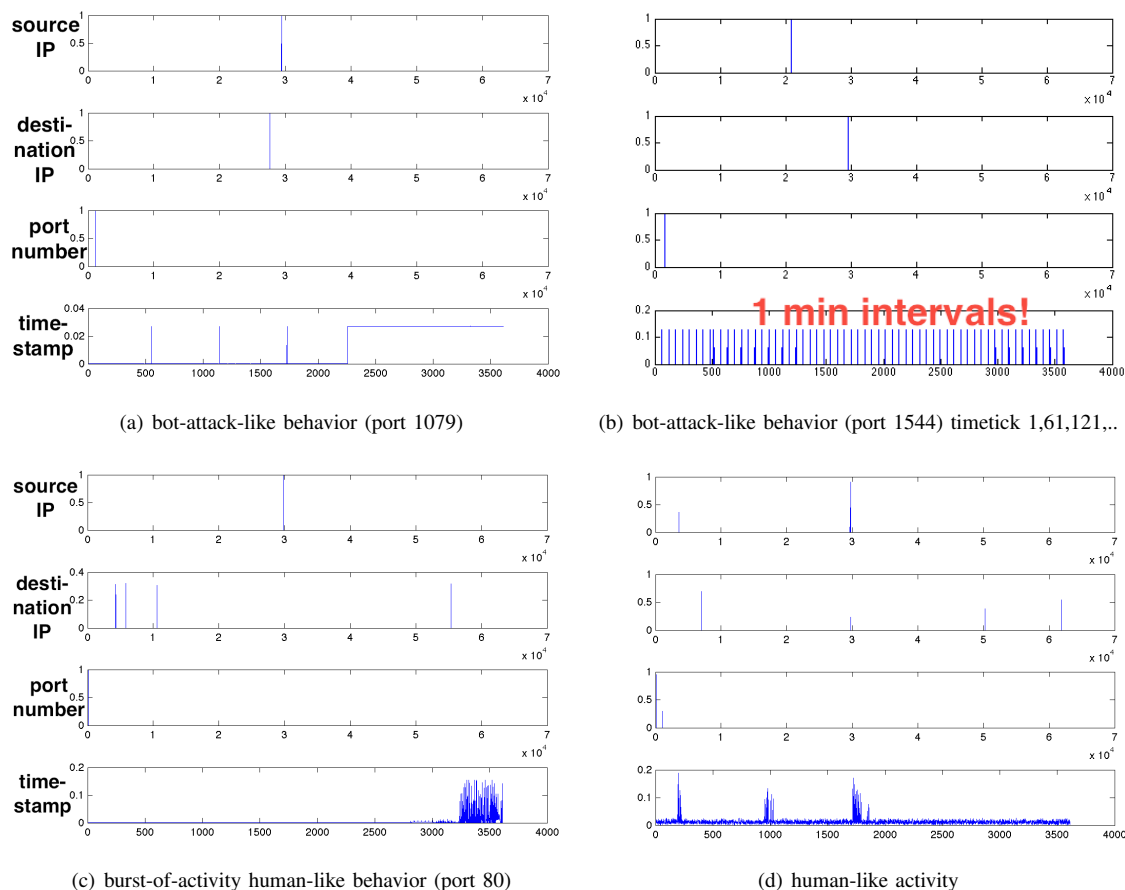


Fig. 4. Anomalies on the LBNL network traffic dataset: (a) The connection is established on port 1079 periodically and eventually it becomes persistent. The first three spikes are nearly evenly spaced, with their time differences being 588, 589, and 525 seconds. (b) Similar bot attack-like behaviour on port 1544 with perfectly evenly spaced spikes of activity, every 60 seconds. (c) An overwhelming burst of traffic on port 80 (HTTP) is concentrated on a confined time interval, which probably indicates a human browsing a website located on the destination IP address. (d) Bursty, human-like traffic

Retrieval (IR) and Artificial Intelligence (AI). In this dataset, each paper is connected to the authors and the venue where it was presented.

- DBLP-2 [8]: The dataset consists of 567,090 papers that appeared in 3,571 different conferences and were written by 418,236 authors over the span of 49 years (1959-2007). The graph from which we extracted the data consists of links between the papers, the conferences where they were presented and their authors.
- LBNL [23]: This dataset is comprised of network traffic measurements, spanning 1 hour, with granularity of seconds. This particular portion of the dataset contains the source IP address, the destination IP address and the port number of all the connections initiated within this hour, along with the timestamp for each connection.
- FACEBOOK [30]: In the FACEBOOK dataset, users are connected to each other based on their posting activity, i.e., there is a link from user  $A$  to user  $B$  if the former posted on the wall of the latter. The data also includes the timestamp of the posting event.

## B. Change detection over time

As we mention in the Introduction, an interesting application is discovering 'bridge'-authors, i.e. researchers that gradually shift from one area to the other. Using the DBLP-1 dataset, we were able to automatically identify a well known professor as a specific example of such 'bridge' author using TENSORSPLAT. In particular, Figure 3 demonstrates the switch of the author from purely Database related conferences to venues whose main focus is Data Mining. In this Figure, we plot the two columns of matrix  $\mathbf{C}$  of the PARAFAC decomposition, which correspond to the *temporal profiles* of the micro-clusters derived from the columns of  $\mathbf{A}$  and  $\mathbf{B}$ . More specifically, each column of  $\mathbf{A}$  will contain 'similar' authors who jointly publish to the set of 'similar' conferences/journals denoted by each column of  $\mathbf{B}$ .

## C. Anomaly detection over time

**LBNL dataset** For this dataset, we formed a 4-way tensor of dimensions  $65K \times 65K \times 65K \times 3.6K$  and applied TENSORSPLAT. Here we report some of the most interesting components of the tensor decomposition.

In Figure 4 we present our findings on the LBNL dataset: Columns of matrix **A** correspond to the source addresses for each micro-cluster, columns of **B** to the destination addresses, columns of **C** to the port number of the connections belonging to the cluster and columns of **D** to the temporal profile of each connection cluster. The first subfigure demonstrates a bot-attack like behavior on port 1079: nearly evenly spaced intervals of activity; the time intervals between the activity spikes have durations of 588, 589, and 525 seconds. The second subfigure, probably our most astounding finding on the LBNL dataset, is also what appears to be a bot-attack like behavior on port 1544, only now, the activity spikes are perfectly evenly distributed. In particular, each spike occurs precisely every 1 minute. Finally, the third subfigure demonstrates a high degree of activity concentration on a small portion of the dataset’s time span. This connection was established on port 80 which is used for HTTP and is probably caused by a human; this behavior possibly indicates a person who is browsing a webpage hosted by a web server located at the destination IP address.

**FACEBOOK dataset** As mentioned in table II, we applied TENSORSPLAT to the graph of  $\sim 64,000$  users who are linked to each other via timestamped wall posts. Again, each column of **A** indicates the Wall owners belonging to one group, each column of **B** singles out the people who post on the Wall(s) of the cluster and each column of **C** shows the temporal activity of the cluster. The case presented in Fig. 5 is a good example of *novelty* detection. We observe that three different people post to the wall of a user on a specific date. This newly discovered behavior may be indicator of a special event (e.g., birthday, job offer, graduation, marriage) of the wall owner.

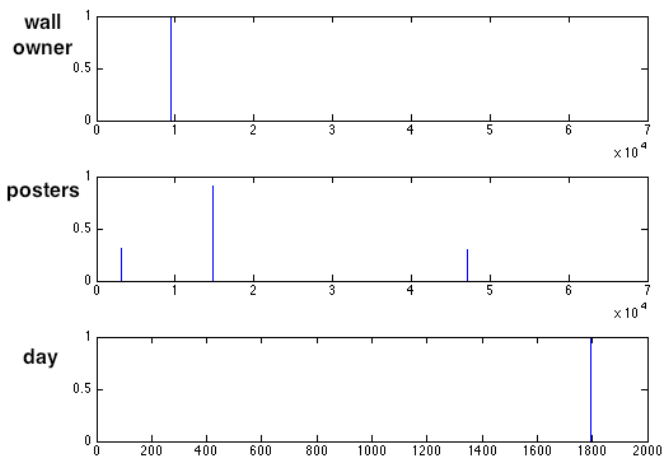


Fig. 5. Birthday-like event: In this plot we illustrate the three vectors of a rank-one triplet, produced by the PARAFAC decomposition. More specifically, we call this emerging pattern ‘Birthday-like event’, because many people (plot/vector 2) are posting on a specific person’s wall (plot/vector 1) on a single date (plot/vector 3). This behavior possibly suggests that there is a special event relevant to the wall owner, for example their birthday or some other type of celebration (e.g. job offer, graduation, marriage etc)

#### D. Clustering

TABLE III  
AUTHORS CLUSTERED BY CONFERENCE

Authors	Conferences
Bing Liu, Jure Leskovec, Christos Faloutsos, Hanghang Tong, Wynne Hsu	KDD
Andrew W. Moore, John Shawe-Taylor, Nello Cristianini, Michael I. Jordan	ICML
Michael J. Carey, H. V. Jagadish, Rakesh Agrawal, Divesh Srivastava, Christos Faloutsos	VLDB
Jeffrey F. Naughton, David J. DeWitt, Nancy E. Hal	SIGMOD
Vincent Conitzer, Tuomas Sandholm, Andrew Gilpin	AAAI

Another interesting application of TENSORSPLAT is the detection of clusters. Specifically, the tensor decomposition gives us nice insights about groups of entities having common one or more of the tensor dimensions. For instance, we may get groups of authors because they publish at the same conferences, at the same time (in this case we identify groups of people who share common research interests), or Like in the motivating example used in the presentation of tensor decomposition, columns of **A** and **B** denote groups of co-authors who are similar in the sense that they publish to (similar) venues, highlighted by the columns of **C**. In Table III we observe some clusters of authors who usually publish to the same conferences. Notice that, in some cases, there is advisor-advisee relationship between the authors that are clustered together.

#### IV. RELATED WORK

**Anomaly Detection** Anomaly detection (or outlier detection, or microcluster detection) involves algorithms that try to spot various types of “strange” behaviors, which deviate “much” from the normal, expected behavior. In general, the problem is not well defined and the anomalousness of an entity heavily depends on the application that is of interest (e.g., malware detection, fraud detection, network intrusion, biological data, image processing etc). Despite the fact that there is no clear definition for anomalies, the anomaly detection consists an important research direction in multitudinous fields. Chandola et al. have written two nice surveys on anomaly detection; the first one [11] refers to the different types of anomalies encountered in different applications, while the second one [12] focuses on outliers in discrete sequences. In a nutshell, the methods that have been used to spot anomalies are classification or clustering-based, statistical, information theoretic or spectral. In the current paper, we focus on a clustering method which accommodates the multiple features (dimensions) of our data.

**Tensor Software** Matlab provides two high quality toolboxes: (a) the Tensor Toolbox [4], [6], which supports operations on sparse tensors, and (b) the N-Way Toolbox [3], which specializes in dense tensors.

**Tensor Applications** Tensors are very powerful tools for the analysis of the continuously increasing (multi-dimensional) data that becomes available, and thus they are rather popular in the data mining field. Among the most successful applications of tensors is the incorporation of textual and topical information in the well known link analysis algorithm HITS [18], as well as the social network analysis of the ENRON dataset in [5]. Other interesting applications of tensors include identification of epileptic seizures [1], improvements of the web search by including information about the clicks [28], network analysis and visualization by a sampling-based tensor decomposition (Tucker3) [27], cross-language information retrieval [14], analysis of large heterogeneous networks [21], as well as image analysis [29]. The applications of tensors are not restricted only in the data mining area; Chemometrics [10] and Signal Processing [26] are only two of the many areas that have successfully used tensors.

## V. CONCLUSIONS

In this paper, we propose TENSORSPLAT, a tensor-based approach for spotting anomalies in graphs. Our method is applicable to a diverse setting of data, highly interpretable (as opposed to other tensor decompositions, e.g. Tucker3), and readily extensible to analysis of multi-way tensors. We demonstrate the practicality of TENSORSPLAT by providing success stories of anomaly and novelty detection, as well as soft clustering of high dimensional data; TENSORSPLAT spots: (i) “bridge” authors, i.e., authors that are publishing in multiple areas, and more ambitiously, authors that changed research area over the time in the large-scale DBLP network, (ii) clusters of authors that often publish together or are connected via advisor-advisee relationship, and (iii) malicious attacks in a large computer network.

So far, our approach requires human intervention at the end, so that the most interesting components of the decomposition are chosen. As future work, we intend to extend the above framework, so that it automatically detects interesting/anomalous patterns without the need for a human to single them out.

## VI. ACKNOWLEDGEMENTS

Funding was provided by the U.S. Army Research Office (ARO), the Defense Threat Reduction Agency under contract No. HDTRA1-10-1-0120 and Defense Advanced Research Projects Agency (DARPA) under Contract Number W911NF-11-C-0088. This work is also partially supported by an IBM Faculty Award. The content of the information in this document does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

## REFERENCES

[1] E. Acar, C. Aykut-Bingol, H. Bingol, R. Bro, and B. Yener. Multiway analysis of epilepsy tensors. *Bioinformatics*, 23(13):i10–i18, 2007.

[2] Leman Akoglu, Mary McGlohon, and Christos Faloutsos. Oddball: Spotting anomalies in weighted graphs. In *PAKDD*, pages 410–421, 2010.

[3] C.A. Andersson and R. Bro. The n-way toolbox for matlab. *Chemometrics and Intelligent Laboratory Systems*, 52(1):1–4, 2000.

[4] Brett W. Bader and Tamara G. Kolda. Efficient MATLAB computations with sparse and factored tensors. *SIAM Journal on Scientific Computing*, 30(1):205–231, December 2007.

[5] B.W. Bader, M.W. Berry, and M. Browne. Discussion tracking in enron email using parafac. *Survey of Text Mining II*, pages 147–163, 2008.

[6] B.W. Bader and T.G. Kolda. Matlab tensor toolbox version 2.2. *Albuquerque, NM, USA: Sandia National Laboratories*, 2007.

[7] Stephen Bay, Krishna Kumaraswamy, Markus G. Anderle, Rohit Kumar, and David M. Steier. Large scale detection of irregularities in accounting data. *ICDM*, 0:75–86, 2006.

[8] The DBLP Computer Science Bibliography. <http://www.informatik.uni-trier.de/~ley/db/>.

[9] Markus M. Breunig, Hans-Peter Kriegel, Raymond T. Ng, and Joerg Sander. LOF: Identifying density-based local outliers. In *SIGMOD*, pages 93–104, Dallas, TX, 2000.

[10] R. Bro. Parafac. tutorial and applications. *Chemometrics and Intelligent Laboratory Systems*, 38(2):149–171, 1997.

[11] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM Comput. Surv.*, 41(3):15:1–15:58, July 2009.

[12] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection for discrete sequences: A survey. *TKDE*, 24:823–839, 2012.

[13] Duen Horng Chau, Carey Nachenberg, Jeffrey Wilhelm, Adam Wright, and Christos Faloutsos. Large scale graph mining and inference for malware detection. In *SDM*, pages 131–142, 2011.

[14] P.A. Chew, B.W. Bader, T.G. Kolda, and A. Abdelali. Cross-language information retrieval using parafac2. In *SIGKDD*, pages 143–152. ACM, 2007.

[15] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman. Indexing by latent semantic analysis. *JASIS*, 41(6):391–407, September 1990.

[16] C. Eckart and G. Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.

[17] R.A. Harshman. Foundations of the parafac procedure: Models and conditions for an “explanatory” multimodal factor analysis. 1970.

[18] T.G. Kolda and B.W. Bader. The tophits model for higher-order web link analysis. In *Workshop on Link Analysis, Counterterrorism and Security*, volume 7, pages 26–29, 2006.

[19] T.G. Kolda and B.W. Bader. Tensor decompositions and applications. *SIAM review*, 51(3), 2009.

[20] D.D. Lee, H.S. Seung, et al. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.

[21] K. Maruhashi, F. Guo, and C. Faloutsos. Multiaspectforensics: Pattern mining on large-scale heterogeneous networks with tensor analysis. In *Advances in Social Network Analysis and Mining*, 2011.

[22] Caleb C. Noble and Diane J. Cook. Graph-based anomaly detection. In *SIGKDD*, pages 631–636, 2003.

[23] R. Pang, M. Allman, M. Bennett, J. Lee, V. Paxson, and B. Tierney. A first look at modern enterprise traffic. In *SIGCOMM Internet Measurement*, pages 2–2. USENIX Association, 2005.

[24] Spiros Papadimitriou, Hiroyuki Kitagawa, Phil Gibbons, and Christos Faloutsos. LOCI: Fast outlier detection using the local correlation integral. *ICDE*, March 2003 5 - March 8 2003.

[25] Karlton Sequeira and Mohammed Javeed Zaki. Admit: anomaly-based data mining for intrusions. In *KDD*, pages 386–395, 2002.

[26] N.D. Sidiropoulos, G.B. Giannakis, and R. Bro. Blind parafac receivers for ds-cdma systems. *TSP*, 48(3):810–823, 2000.

[27] J. Sun, S. Papadimitriou, C.Y. Lin, N. Cao, S. Liu, and W. Qian. Multivis: Content-based social network exploration through multi-way visual analysis. In *SDM*, volume 9, pages 1063–1074, 2009.

[28] J.T. Sun, H.J. Zeng, H. Liu, Y. Lu, and Z. Chen. Cubesvd: a novel approach to personalized web search. In *WWW*, pages 382–390. ACM, 2005.

[29] M. Vasilescu and D. Terzopoulos. Multilinear analysis of image ensembles: Tensorfaces. *ECCV*, pages 447–460, 2002.

[30] Bimal Viswanath, Alan Mislove, Meeyoung Cha, and Krishna P. Gummadi. On the evolution of user interaction in facebook. In *WOSN*, August 2009.