Tensor Clustering: A Review

Georgios Drakopoulos
Department of Informatics
Ionian University
c16drak@ionio.gr

Evaggelos Spyrou
NCSR "Demokritos"
espyrou@iit.demokritos.gr

Phivos Mylonas
Department of Informatics
Ionian University
fmylonas@ionio.gr

***Abstract*—Tensor algebra is the next evolutionary step of linear algebra to more than two dimensions. Its plethora of applications include signal processing, big data, deep learning, multivariate numerical analysis, information retrieval, and social media analysis. As is precisely the case with data matrices, decompositions and factorizations with special properties reveal inherent but latent patterns which are not immediately discernible. Alternatively, for large tensors direct clustering can yield similar patterns. Once identified, said patterns can pave the way for other operations commonly found in a knowledge mining pipeline such as compression, outlier discovery, and higher order statistics. This survey concisely presents the key tensor clustering techniques as well as their applications. Additionally, deep learning frameworks which natively support tensors such as TensorFlow, Breeze, Spark MLlib, and Tensor Toolbox are presented.**

***Index Terms*—tensor algebra; multilayer graphs; knowledge mining; tensor stack network; multilinear discriminant function; TensorFlow; theano; keras; MLlib; Breeze**

1. Introduction

In the 6V era, tensor algebra currently garners considerable multidisciplinary attention. Two major drivers behind this intense interest is the multitude of close ties to other resarch fields which heavily rely on multivariate linear algebra as well as the fact that tensor algebra has more expressive power. Thus, seemingly complex formulas involving multiple sums, such as Volterra or multivariate Taylor expansions, can be rewritten in a very natural way, allowing thus for almost intuitive interpretability and for deeper understanding. Therefore, simultaneous linear interaction can be captured in applications so diverse such as brain circuit simulation, numerical deep learning, information retrieval, supply chain and logistics networks, non-linear system identification, computational combinatorics, multispectral image processing, and multilayer graph mining.

Besides the algorithmic cornerstones, efficient software tools having tensors as a native data type are required. The latter have already been implemented in a number of MATLAB toolboxes such as Tensor Toolbox and TensorLab. TensorFlow in an open source low level platform with C++ and Python APIs originally developed by Google whose primary data unit is, as its name suggests, a tensor. Low level deep learning operations over GPU are also implemented in theano, whereas keras provides high level front ends for both TensorFlow and theano. DistKeras is a recent keras variant which can be executed over Spark. Tensor operations are being currently integrated in MLlib for Spark and at a higher level at Breeze for Scala. Additionally, parallel linear algebra libraries are being modified for CUDA.

The primary research objective of this survey is a broad review of tensor clustering techniques. Where appropriate, the fundamental tensor algebra background is provided for deeper understanding. As a secondary objective the mainstream tensor related software is reviewed.

The structure of this work follows. Section 2 describes milestones in tensor research, tensor related scientific software, and outlines the frame of reference. Tensor clustering criteria are discussed in section 3. Important applications are mentioned in section 4. Finally, section 5 describes future research directions.

2. Previous Work

2.1. Section Overview

This section focuses on scientific literature regarding tensors and relevant topics such as higher order signal processing, which forms a starting point for the proposed research. Moreover, important software as well as a number of digital repositories are mentioned. Finally, for completeness purposes, certain topics related to tensor algebra are mentioned even though they will not be part of the proposed research.

At this point it should be emphasized that since 2013 in Google scholar the number of tensor related publications has been increasing almost exponentially, with contributions coming from every major US and EU university. Additonally, the same trend can be traced at the number of special sessions in deep learning and data mining conferences. These scientometric data indicate the current intense research interest for tensors.

978-1-7281-3634-9/19/$31.00

2.2. Tensor Algebra

Tensor algebra is a superset of matrix algebra [1] [2] since tensors are multidimensional vectors indexed by a tuple of p integers in the same way a matrix is a two dimensional vector [3] [4] [5]. The importance of tensors lies in their greater expressive power and the wider ranger of operations which can be defined on them compared to matrices [6] [7]. Kruskal decomposition [8] as well as Tucker [9] [10] and Poisson factorizations [11] are among the common tensor operations.

2.3. Higher Order Signal Processing

Higher order signal processing extended classical signal processing to more than two input variables and to the study of correlations of three or more points in time or in frequency. Such methods are necessary in important scenaria such as nonlinear system analysis, cyclostationary signal analysis, blind source separation. A typical example of the latter is ICA [12] [13], which under mild assumptions blindly separates a set of signals driven to a system given only the output that system. Typical higher order statistics include higher order moments [14] [15] and cumulants [16] [17] [18].

2.4. Multilayer Graphs

Multilayer graphs extend ordinary graphs by assigning to each vertex a not necessarily distinct label and by allowing the existence of multiple edges between the same vertex pair as long as the edges have pairwise distinct labels. One of the algebraic counterparts of this combinatorial object is a third order tensor with a specific entries, appropriately called an *adjacency tensor*. Note that the converse is not always true, as not every third order tensor corresponds to a multilayer graph, in the same way that not every matrix is a graph adjacency matrix [19] [20]. Thus, multilayer graphs represent a restricted class of third order tensors [21].

2.5. Tensor Software

Due to the intense research on tensors there is a significan amount of high quality software. MATLAB toolboxes include Tensor Toolbox and N-way array toolbox, which focus on multilinear algebra, and TensorLab which is oriented towards higher order signal processing. Breeze is a linear algebra library for Java and Scala inherently supporting multidimensional vectors. Concerning deep learning, open source TensorFlow from Google employs tensors as the basic data unit, while GraphLab can be modified to perform tensor operations. Finally, NumPy and SciPy offer tensor functionality in Python.

Multilayer graphs constitute the basic building blocks in graph databases, which are mentioned separately in the next subsection. Also, NetworkX Python library natively supports multilayer graphs, provided their size can fit in the main memory.

Because of the potentially very large size and their very low density, which in certain applications can be as low as 5-7%, there is an intense research interest for distributed tensor processing in systems such as Hadoop, Spark, or Flink. The difference between the distributed computation model in these systems is highly likely to lead to different implementations of the same tensor operations. Factors influencing these implementations are memory requirements, the number of available computational nodes, fault tolerance, and scheduling both at the system and the node level. In any case, the application programmer maintains a significant degree of flexibility.

Similar functionality, with different operational characteristics, is offered by CUDA in GPUs. Specifically, GPU parallel computation compared to that of a distributed computation is quicker, more local, and more reliabile. However, this comes at the cost of lower system memory and reduced programming flexibility, again in comparison to distributed computing [22] [23].

2.6. NoSQL Databases

The recent advent of NoSQL databases created new possibilities for handling structured data in formats other than tabular. Moreover, they can store semistructured or even unstructured data [24] [25]. The NoSQL ecosystem consists of four primary technologies listed in table 1.

Graph databases are ideal for representing and managing multilayered graphs and, thus, indirectly adjacency tensors. In fact, most commercially available graph databased support multigraphs inherently and can be used as a platform for developing multilayer graph analytics such as centrality and cohesion metrics or community structure discovery algorithms.

2.7. Online Repositories

It is common knowledge in deep learning and data science that data factors such as volume, veracity, and variability are closely liked to clustering effectiveness, probably in a critial manner in certain domains. Therefore, digital repositories whose datasets are of high quality, according to common metrics, concerning those factors are of paramount importance. Online archives whose datasets are often employed as benchmarks for assessing clustering behavior include Kaggle, Georgia Tech GIS Center, Stanford Network Analysis Project (SNAP), UFL Matrix Collection, and UCR Dataset Collection.

2.8. Related Fields

Tensors and multidimensional arrays are also present under various names and forms in other fields of computer science and engineering not mentioned above. In the world of relational databases, OLAP cubes are three or four dimensional arrays which may contain categorical data or strings [26] [27].

TABLE 1. NoSQL DATABASES

Type	Data type	Standard	Software
Key-value	Associative array	JSON, XML	Apache Dynamo, Riak, Redis, Oracle NoSQL
Column family	Long rows	JSON, BSON	Apache Cassandra, Keyspace
Document	Documents	JSON, BSON, YAML, XML	MongoDB, CouchDB, OrientDB, CreateIO
Graph	Linked data	RDF, JSON-LD, conceptual multigraph	Neo4j, TitanDB, Sparksee, InfiniteGraph

In [28] is described a versatile, efficient, and effective linear algebraic technique for converting first order digital influence metrics in social media to higher order ones. The latter are more appropriate for discovering structural or functional patterns of interest in graphs, since by definition graphs rely on recursive link, communication, and functionality diffusion in general between their vertices. This is the reason that PageRank is by far more successful than previous vertex centrality metrics such as degree centrality.

A heuristic for tensor clustering along with the associated algorithmic details is presented in [29]. Specifically, a genetic algorithm for the approximate clustering of a third order tensor containing geolocational and linguistic Twitter data coming from the *de jure* and *de facto* trilingual grand duchy of Luxembourg is described. Two alternative objective functions have been used, each equally based on linguistic change models and a Zipf function of geodesical distance, whereas the selection, crossover, and mutation, operators have been linked to community merging and splitting.

Finally, [30] revolves around a conceptual framework for ontologies with labeled edges. Thus, where the underlying domain permits it, mutiple connections can exist between two given entities. As a concrete example, a selected number of persons from the 1970s and 1980s Apple and the connections between them have been extracted from the official biography of Steve Jobs and the 1999 film *The pirates of Silicon Valley*.

3. Clustering Critria

Depending on the underlying domain, clustering can take a different form in order to optimize execution time, memory, or interpretability, the latter being of paramount importance as it is inherently tied to clustering validity. For instance, if the tensor contains fMRI images, then a neurophysiological interpretaion based on a commonly acceptable brain atlas such as AAL or AAL2 is necessary. Similarly, if the tensor comprises of social media account interactions, then clustering results should accord with the principles of digital activity found in the scientific literature.

The following complexity criteria typically act as a primary benchmark for tensor clustering:

- **C1**: The total execution time as well as any critical paths in execution flow.
- **C2**: The total demands across the memory hierarchy as well as access patterns.
- **C3**: The scalability expressed as the ability of the algorithm to be implemented in a GPU card or in a distributed platform. In the first case the fundamantal restrictions are the card memory and the communication protocol between the card processors, whereas in the second case the node communication protocols and the full exploitation of the distributed computation model, e.g. DAG in Spark and MapReduce in Hadoop, are paramount factors.

Secondary criteria for evaluating the clustering algorithm besides the complexity related ones are listed below. They represent other significant viewpoints and are ordered from the most general to the most specific.

- **C4**: The evaluation of the role of the entire system in the algorithm performance including architecture, database type, scalability, and implementation language.
- **C5**: Domain specific as well as generic quality clustering metrics to ensure comparison fairness.
- **C6**: Memory effect, in the form of the architecture of possibly nested feedback loops, in deep learning scenaria where convolutional and feedforward networks are involved as units of tensor networks.
- **C7**: Data coding effect, especially for heuristics such as tensor stack networks and tensor genetic algorithms.
- **C8**: Data distribution effect in probabilistic scenaria with emphasis in the interpretability of both input and output distribution. For instance, in social network scenaria it might be of interest that the message frequency distribution has a Poisson distribution or that hashtag length has a Zipf distribution.
- **C9**: Numerical stability for a large number of floating point operations, which is an important factor in computing probabilities or solving large linear systems, operations which lie at the kernel of many deep learning problems. Relevant considerations are the selection of floating point system and the possible benefits if using numbers which are 64- and 128-bit long, provided that specialized software exists.

4. Applications

This section focuses on tensor applications. Given the topic diversity and the length of this section, it is evident that tensors have a very broad spectrum of applications in computer science and engineering. This happens because of the following reasons:

- Being the next evolutionary step of matrix algebra, tensor alebra can appear virtually everywhere when the former can, with certain necessarynotation rearrangements. Perhaphs the most representative example is the the celebrated singular value decomposition (SVD) or latent semantic indexing (LSI), as it is known in information retrieval. The basis of SVD is the factorization of a, potentially large, data matrix in three simpler matrices with special properties, which can be more naturally expressed as two tensor products between these factors.

- Moreover, tensors almost by definition appear in multivariable Taylor and McLaurin expansions. Under the viewpoint of Newton's work, who proved that any arbitrary function is essentially a polynomical of possibly very high degree, tensors can naturally represent or approximate any nonlinear system, which is not possible in the general case in matrix algebra.

Data mining is a broad field encompassing a multitude of diverse techniques with matrix algebra constitiing one of the maintays of the field. Tensors can be the primary data represenatation, especially when the input data are images or matrices. Additionally, tensors can appear at the preprocessing stage of the general data mining piepline. Moreover, tensors have been used as a starting point of the vector clustering algorithm proposed in [31], where a graph is indirectly constructed in order to provide an initial estimate of the cluster number for the original version of the k-means algorithm.

Online social networks have emerged as the primary theater of the connected era for digital marketing, cultural content dissemination, and political campaigns among others. Ultimately, they constitute a vehicle for massive and multidimensional communication, as in Turkey in the wake of the failed military coup of 2016 and during the Arab Spring of 2011, or for the outcome prediction of political and social events. However, as clearly indicated by the recent experience with the US and French presidential elections, more reliable prediction techniques should be invented.

Tensors can represent a social network in higher granularity compared to the classical adjacency matrix representation, as adjacency tensors can take into account separately each of the factors associated with netizen interaction including geographical and linguistic criteria [29] [32]. Furthermore, each of these factors can be the building block for digital influence metric and, in turn, those metrics can be fused through tensor algebra to advanced and higher order influence metrics which yield deeper insight into information diffusion as in [33] and [28]. Digital influence is a fundamental question with applications ranging from fake news discovery to digital marketing. As a general note, these factors can be determined either by sociopolitical critreria [34] [35] or by the interaction options offered by the social network itself [36].

Community models based on tensor algebra [37] [38] address the important aspect of community structure discovery, an inherent property of the large scale free graphs which can be found literally everywhere in the real world.

Some recent advances is cognitive science and in neuroinformatics can be attributed in part to the successful algorithmic analysis of large brain related data such as CAT scans, fMRI video, and EEG time series. Tensor algebra has already been applied to fMRI with positive results [39] [40], although there is still room for improvement, as the set of features may vary in each clinical case or for each type of diagnosis. Given brain complexity, both in terms of connectivity and functionality, it is highly unlikely that a single tensor analysis suffice because of the inherent multidinensionality of brain data [41] [42]. However, it is reasonable to expect that tensor based brain models which take into account variables such as time, space, subject condition, and subject task are a mathematically sound starting point for brain study [43].

Traditionally ontologies are represented as directed trees such as the XML trees which are heavily used at the Semantic Web. The edges of these ontological trees have attributes denoting predicates. Multilayer graphs can also be used in this context and they can replace complex relationships between two given entities with a number of simpler ones. Thus, large relationship sets can be reduced to smaller and simpler ones, resulting in lower complexity and added parsimonity. Moreover, multilayer graphs can represent time varying or incomplete ontologies by storing mutiple instances of the same ontology using techniques such as those proposed in [44].

Tensors can be used in information retrieval in order to extend in a number of ways the classical term-document matrix model. Such extentions include for instance the term-author-document and the term-keyword-document third order tensors [45]. Besides their accordance with intuition, such tensor models can take advantage of clustering algorithms tailored for three dimensions instead of the generic schemes.

In deep learning tensors take many forms including linear tensor discriminant analysis (LTDA) [46] and tensor stack networks (TSN) [47]. Both schemes have been successfully applied to many challenging tasks including speech recognition [48], face recognition [49], gait classification [50], and automated object identification in computer vision [51]. These methodologies require a relatively large number of data in order to reliably construct an indirect partition of the data space. Thus, new schemes or improved variations of existing ones can be derived.

Besides the abovementioned applications, tensors in one form or another can be found in other fields. Even if this subsection is by no means exhaustive, its breadth allows the reader to understand the wide applicability of tensor algebra in computer science and engineering.

In computational biology multilayer graphs can represent protein-protein interactions or time varying genomic data [52].

In transportation networks and logistics studies tensors can represent multiple connections such as road and train networks or alternative routes offered by various logistics companies, allowing a detailed investigation of supply and distribution chains. The same hold true for seaways and airways, as different companies may offer pricing policies depending on the conditions in different geographical areas.

Nonlinear system identification can be reduced through expansions such as Taylor and Volterra series to the numerically stable determination of the coefficients of multilinear and multivariable models [53] [54].

Multispectral image analysis relies on the separate study and manipulation of the spectra of the same image corresponding to the various depictable wavelengths [55]. Applications include 3D reconstruction of damaged skin areas [56] and the noninvasive burn depth determination in preparation for restoration operations [57].

Finally, in scientific computing, a number of transforms can exploit an alternative tensor representation of data in order to achieve a considetable speedup [58] [59].

5. Conclusions

This survey reviews methods for tensor clustering as well as scientific software for efficiently performing data intensive tensor operations. Tensor clustering plays an instrumental role in a wide array of fields, most prominently deep learning, ontology, fMRI analysis, big data management, information retrieval, non-linear system identification, and knowledge disovery. Fundamental tensor algebra functionality is already provided by both deep learning software such as TensorFlow and theano as well as by libraries for linear algebra such as Tensor Toolbox and TensorLab for MATLAB and Breeze for Scala.

Possible research directions include the developement of specialized efficient operations for lower order tensors, e.g. for tensors with up to five dimensions. Symmetries drawn from group theory can be used to accelerate various operations such as clustering or Kruskal factorization. Additionally, at a lower level, the recent advances in GPU computing and in distributed systems should be exploited by appropriate layers of tensor related software.

Acknowledgment

This research has been co-financed by the European Union and Greek national funds through the Competitiveness, Entrepreneurship and Innovation Operational Programme, under the Call "Research - Create - Innovate", project title: "Development of technologies and methods for cultural inventory data interoperability", project code: T1EDK-01728, MIS code: 5030954.

References

[1] L. De Lathauwer, B. De Moor, and J. Vandewalle, "Independent component analysis based on higher-order statistics only," in *Proceedings of the 8th IEEE Signal Processing Workshop on Statistical Signal and Array Processing*. IEEE, 1996, pp. 356–359.

[2] L. De Lathauwer and J. Vandewalle, "Dimensionality reduction in higher-order signal processing and rank-$(r_1, r_2, \ldots, r_n)$ reduction in multilinear algebra," *LAA*, vol. 391, pp. 31–55, 2004.

[3] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM Review*, vol. 51, no. 3, pp. 455–500, 2009.

[4] T. G. Kolda, "Orthogonal tensor decompositions," *SIAM Journal on Matrix Analysis and Applications*, vol. 23, no. 1, pp. 243–255, 2001.

[5] J. B. Kruskal, "More factors than subjects, tests and treatments: An indeterminacy theorem for canonical decomposition and individual differences scaling," *Psychometrika*, vol. 41, no. 3, pp. 281–293, 1976.

[6] M. Filipović and A. Jukić, "Tucker factorization with missing data with application to low-n-rank tensor completion," *Multidimensional systems and signal processing*, vol. 26, no. 3, pp. 677–692, 2015.

[7] D. V. Savostyanov, E. E. Tyrtyshnikov, and N. L. Zamarashkin, "Fast truncation of mode ranks for bilinear tensor operations," *LAA*, vol. 19, no. 1, pp. 103–111, 2012.

[8] B. W. Bader and T. G. Kolda, "Efficient MATLAB computations with sparse and factored tensors," *SIAM Journal on Scientific Computing*, vol. 30, no. 1, pp. 205–231, 2007.

[9] N. Lee and A. Cichocki, "Fundamental tensor operations for large-scale data analysis using tensor network formats," *Multidimensional Systems and Signal Processing*, pp. 1–40, 2017.

[10] I. Perros, R. Chen, R. Vuduc, and J. Sun, "Sparse hierarchical Tucker factorization and its application to healthcare," in *ICDM*. IEEE, 2015, pp. 943–948.

[11] P. Gopalan, J. M. Hofman, and D. M. Blei, "Scalable recommendation with Poisson factorization," *arXiv preprint:1311.1704*, 2013.

[12] P. Comon, J. M. Ten Berge, L. De Lathauwer, and J. Castaing, "Generic and typical ranks of multi-way arrays," *LAA*, vol. 430, no. 11, pp. 2997–3007, 2009.

[13] L. De Lathauwer, P. Comon, B. De Moor, and J. Vandewalle, "ICA algorithms for 3 sources and 2 sensors," in *Proceedings of the IEEE Signal Processing Workshop on Higher-Order Statistics*. IEEE, 1999, pp. 116–120.

[14] R. Pan and C. L. Nikias, "The complex cepstrum of higher order cumulants and nonminimum phase system identification," *IEEE transactions on acoustics, speech, and signal processing*, vol. 36, no. 2, pp. 186–205, 1988.

[15] C. L. Nikias and J. M. Mendel, "Signal processing with higher-order spectra," *IEEE Signal processing magazine*, vol. 10, no. 3, pp. 10–37, 1993.

[16] G. Agarwal and A. Biswas, "Inseparability inequalities for higher order moments for bipartite systems," *New Journal of Physics*, vol. 7, no. 1, p. 211, 2005.

[17] A. N. Shiryaev, "Some problems in the spectral theory of higher-order moments. I," *Theory of Probability & Its Applications*, vol. 5, no. 3, pp. 265–284, 1960.

[18] V. Leonov and A. N. Shiryaev, "Some problems in the spectral theory of higher-order moments. II," *Theory of Probability & Its Applications*, vol. 5, no. 4, pp. 417–421, 1960.

[19] Y. Yang, S. Han, T. Wang, W. Tao, and X.-C. Tai, "Multilayer graph cuts based unsupervised color-texture image segmentation using multivariate mixed Student's t-distribution and regional credibility merging," *Pattern recognition*, vol. 46, no. 4, pp. 1101–1124, 2013.

[20] M. De Domenico, V. Nicosia, A. Arenas, and V. Latora, "Structural reducibility of multilayer networks," *Nature communications*, vol. 6, 2015.

[21] G. Drakopoulos, P. Gourgaris, and A. Kanavos, "Graph communities in Neo4j: Four algorithms at work," *Evolving Systems*, June 2018.

[22] G. Drakopoulos, X. Liapakis, G. Tzimas, and P. Mylonas, "A graph resilience metric based on paths: Higher order analytics with GPU," in *ICTAI*. IEEE, November 2018.

[23] G. Drakopoulos, M. Marountas, X. Liapakis, G. Tzimas, P. Mylonas, and S. Sioutas, "Blockchain for mobile health applications: Acceleration with GPU computing," in *GeNeDis 2018*, P. Vlamos, Ed. Springer, 2018.

[24] J. Webber, "A programmatic introduction to Neo4j," in *Proceedings of the 3rd annual conference on Systems, programming, and applications: Software for humanity*. ACM, 2012, pp. 217–218.

[25] M. Stonebraker, "SQL databases v. NoSQL databases," *Communications of the ACM*, vol. 53, no. 4, pp. 10–11, 2010.

[26] C. Ordonez and Z. Chen, "Evaluating statistical tests on OLAP cubes to compare degree of disease," *IEEE Transactions on Information Technology in Biomedicine*, vol. 13, no. 5, pp. 756–765, 2009.

[27] S. Vinnik and F. Mansmann, "From analysis to interactive exploration: Building visual hierarchies from OLAP cubes," in *EDBT*. Springer, 2006, pp. 496–514.

[28] G. Drakopoulos, A. Kanavos, P. Mylonas, and S. Sioutas, "Defining and evaluating Twitter influence metrics: A higher order approach in Neo4j," *SNAM*, 2017.

[29] G. Drakopoulos, F. Stathopoulou, G. Tzimas, M. Paraskevas, P. Mylonas, and S. Sioutas, "A genetic algorithm for discovering linguistic communities in spatiosocial tensors with an application to trilingual Luxembourg," in *MHDW*, June 2017.

[30] G. Drakopoulos, A. Kanavos, D. Tsolis, P. Mylonas, and S. Sioutas, "Towards a framework for tensor ontologies over Neo4j: Representations and operations," in *IISA*, August 2017.

[31] G. Drakopoulos, P. Gourgaris, A. Kanavos, and C. Makris, "A fuzzy graph framework for initializing k-Means," *IJAIT*, vol. 25, no. 6, pp. 1–21, 2016.

[32] G. Drakopoulos, F. Stathopoulou, A. Kanavos, M. Paraskevas, G. Tzimas, P. Mylonas, and L. Iliadis, "A genetic algorithm for spatiosocial tensor clustering: Exploiting TensorFlow potential," *Evolving Systems*, January 2019.

[33] G. Drakopoulos, "Tensor fusion of social structural and functional analytics over Neo4j," in *IISA*. IEEE, July 2016.

[34] M. D. Conover, B. Gonçalves, J. Ratkiewicz, A. Flammini, and F. Menczer, "Predicting the political alignment of Twitter users," in *SocialCom*. IEEE, 2011, pp. 192–199.

[35] Y.-R. Lin, J. Sun, P. Castro, R. Konuru, H. Sundaram, and A. Kelliher, "Metafac: Community discovery via relational hypergraph factorization," in *SIGKDD*. ACM, 2009, pp. 527–536.

[36] L. A. Adamic and N. Glance, "The political blogosphere and the 2004 US election: Divided they blog," in *Proceedings of the 3rd international workshop on Link discovery*. ACM, 2005, pp. 36–43.

[37] N. Durak, T. G. Kolda, A. Pinar, and C. Seshadhri, "A scalable null model for directed graphs matching all degree distributions: In, out, and reciprocal," in *2nd NSW*. IEEE, 2013, pp. 23–30.

[38] C. Seshadhri, A. Pinar, and T. G. Kolda, "Triadic measures on graphs: The power of wedge sampling," in *ICDM*. SIAM, 2013, pp. 10–18.

[39] D. Werring, C. Clark, G. Parker, D. Miller, A. Thompson, and G. Barker, "A direct demonstration of both structure and function in the visual system: Combining diffusion tensor imaging with functional magnetic resonance imaging," *Neuroimage*, vol. 9, no. 3, pp. 352–361, 1999.

[40] I. Davidson, S. Gilpin, O. Carmichael, and P. Walker, "Network discovery via constrained tensor analysis of fMRI data," in *SIGKDD*, 2013, pp. 194–202.

[41] J.-M. Papy, L. De Lathauwer, and S. Van Huffel, "Exponential data fitting using multilinear algebra: The single-channel and multi-channel case," *Numerical linear algebra with applications*, vol. 12, no. 8, pp. 809–826, 2005.

[42] ——, "Exponential data fitting using multilinear algebra: The decimative case," *Journal of Chemometrics*, vol. 23, no. 7-8, pp. 341–351, 2009.

[43] A. Kachenoura, L. Albera, L. Senhadji, and P. Comon, "ICA: A potential tool for BCI systems," *IEEE Signal Processing Magazine*, vol. 25, no. 1, pp. 57–68, 2008.

[44] S. Kontopoulos and G. Drakopoulos, "A space efficient scheme for graph representation," in *ICTAI*. IEEE, November 2014.

[45] G. Drakopoulos, A. Kanavos, I. Karydis, S. Sioutas, and A. G. Vrahatis, "Tensor-based semantically-aware topic clustering of biomedical documents," *Computation*, vol. 5, no. 3, May 2017.

[46] G. Baudat and F. Anouar, "Generalized discriminant analysis using a kernel approach," *Neural computation*, vol. 12, no. 10, pp. 2385–2404, 2000.

[47] B. Hutchinson, L. Deng, and D. Yu, "Tensor deep stacking networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1944–1957, 2013.

[48] D. Yu, L. Deng, and F. Seide, "The deep tensor neural network with applications to large vocabulary speech recognition," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 21, no. 2, pp. 388–396, 2013.

[49] S. Yan, D. Xu, Q. Yang, L. Zhang, X. Tang, and H.-J. Zhang, "Multilinear discriminant analysis for face recognition," *IEEE Transactions on Image Processing*, vol. 16, no. 1, pp. 212–220, 2007.

[50] D. Tao, X. Li, X. Wu, and S. J. Maybank, "General tensor discriminant analysis and Gabor features for gait recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 10, 2007.

[51] H. Lu, K. N. Plataniotis, and A. N. Venetsanopoulos, "Uncorrelated multilinear discriminant analysis with regularization and aggregation for tensor object recognition," *IEEE Transactions on Neural Networks*, vol. 20, no. 1, pp. 103–123, 2009.

[52] G. Drakopoulos, S. Kontopoulos, and C. Makris, "Eventually consistent cardinality estimation with applications in biodata mining," in *SAC 2016*. ACM, April 2016.

[53] G. Favier and T. Bouilloc, "Parametric complexity reduction of Volterra models using tensor decompositions," in *EUSIPCO*. IEEE, 2009, pp. 2288–2292.

[54] T. Bouilloc and G. Favier, "Nonlinear channel modeling and identification using baseband Volterra–PARAFAC models," *Signal Processing*, vol. 92, no. 6, pp. 1492–1498, 2012.

[55] J.-I. Park, M.-H. Lee, M. D. Grossberg, and S. K. Nayar, "Multispectral imaging using multiplexed illumination," in *ICCV*. IEEE, 2007, pp. 1–8.

[56] C. Kuo, O. Coquoz, T. L. Troy, H. Xu, and B. W. Rice, "Three-dimensional reconstruction of in vivo bioluminescent sources based on multispectral imaging," *Journal of biomedical optics*, vol. 12, no. 2, pp. 024 007–024 007, 2007.

[57] W. Eisenbeiß, J. Marotz, and J.-P. Schrade, "Reflection-optical multispectral imaging method for objective determination of burn depth," *Burns*, vol. 25, no. 8, pp. 697–704, 1999.

[58] K. L. Mickus and J. H. Hinojosa, "The complete gravity gradient tensor derived from the vertical component of gravity: A Fourier transform technique," *Journal of Applied Geophysics*, vol. 46, no. 3, pp. 159–174, 2001.

[59] J. Wess and B. Zumino, "A Lagrangian model invariant under supergauge transformations," *Physics Letters B*, vol. 49, no. 1, pp. 52–54, 1974.