REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
21-06-2018	Final Report	1-Jul-2013 - 28-Aug-2016

4. TITLE AND SUBTITLE
Final Report: YIP: Learning, Dynamics and Intervention in Large-Scale Social Networks

5a. CONTRACT NUMBER
W911NF-13-1-0084

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
611102

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

6. AUTHORS

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES
University of California - Irvine
141 Innovation Drive, Suite 250

Irvine, CA 92697 -7600

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)
U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S)
ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
63758-NS-YIP.5

12. DISTRIBUTION AVAILIBILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Animashree Anandkumar
UU	UU	UU	UU		19b. TELEPHONE NUMBER
					949-824-9072

Standard Form 298 (Rev 8/98)
Prescribed by ANSI Std. Z39.18

RPPR Final Report
as of 09-Jul-2018

Agency Code:

Proposal Number: 63758NSYIP **Agreement Number: W911NF-13-1-0084**

INVESTIGATOR(S):

Name: Animashree Anandkumar
Email: a.anandkumar@uci.edu
Phone Number: 9498249072
Principal: Y

Organization: **University of California - Irvine**
Address: 141 Innovation Drive, Suite 250, Irvine, CA 926977600
Country: USA
DUNS Number: 046705849 EIN: 952226406

Report Date: 30-Nov-2016 Date Received: 21-Jun-2018
Final Report for Period Beginning 01-Jul-2013 and Ending 28-Aug-2016
Title: YIP: Learning, Dynamics and Intervention in Large-Scale Social Networks
Begin Performance Period: 01-Jul-2013 **End Performance Period:** 28-Aug-2016
Report Term: 0-Other
Submitted By: Animashree Anandkumar Email: a.anandkumar@uci.edu
Phone: (949) 824-9072
Distribution Statement: 1-Approved for public release; distribution is unlimited.

STEM Degrees: 3 **STEM Participants:**

Major Goals: *Learning in high dimensions
*Practical algorithms with strong guarantees
*Method of moments at scale through tensor algebra

Accomplishments: pdf attached

Training Opportunities: 3 graduate students were trained as part of this project. They went on to receive their PhD.

Results Dissemination: Numerous talks were given in conferences, workshops and other educational venues. Podcasts and online videos resulted in widespread dissemination

Honors and Awards: The PI received Microsoft faculty award, NSF career award, Google faculty award. Recently, the PI became a named professor at Caltech, the highest academic honor at Caltech

Protocol Activity Status:

Technology Transfer: We have released many open source software packages based on the tensor algorithms developed as part of this project. We have contributed to Tensorly package \url{https://github.com/tensorly} and actively developing it. The tensor decomposition algorithm for topic modeling is now deployed at AWS SageMaker platform and is part of topic modeling framework in AWS Comprehend, the NLP service.

PARTICIPANTS:

Participant Type: PD/PI
Participant: Animashree Anandkumar
Person Months Worked: 12.00 **Funding Support:**
Project Contribution:
International Collaboration:
International Travel:
National Academy Member: N
Other Collaborators:

RPPR Final Report
as of 09-Jul-2018

1 Science

1. What is the mathematical objective of your project? What question are you trying to answer?

   Today we are facing a "data deluge" in almost every domain. Online social networks have seen an explosion in activity and have fundamentally transformed the nature of human interaction. In the biological realm, modern genome sequencers can output data at a rate 400 times faster than the ones a decade ago, and so on. However, although having a transformative potential, the data deluge has not yet been exploited to the fullest extent. Ironically, the data deluge has also resulted in a "data desert". The collected data in many domains are noisy, subsampled, with typically a large number of variables or "unknowns" compared to the number of observations or the "knowns". Such high-dimensionality entails practical principled approaches for learning from ill-posed and ill-behaved data.

   Some of the fundamental questions in high-dimensional learning are: Can we design **scalable** models for efficiently representing and learning high-dimensional data? Here, scalability refers to low **computational** requirements and reduced **sampling** of high-dimensional data. Not all phenomena can be learnt in a scalable manner. Can we characterize the **fundamental limits** on complexity of learning complex phenomena?

2. What are the challenges in doing this? What makes it difficult?

   Learning in high dimensional regime is an ill-posed problem. It is akin to finding a needle in the haystack. In the worst case, learning entails exponential requirements in the amount of samples and computation. Most machine learning tasks require solving non-convex optimization problems. For these problems, as the data dimensions grow, the number of critical points can grow exponentially. Local search methods such as gradient descent can get stuck in one of these critical points. Finding the globally optimal solution is computationally hard in the worst case for non-convex optimization.

3. What is the scientific opportunity that is enabling you to make progress in this difficult area?

   Instead, over the last few years, focus has shifted in characterizing transparent conditions for non-convex problems which are tractable. My research shows that in many instances, these conditions turn out to be mild and natural for machine learning applications. I demonstrate that a range of learning problems can be solved efficiently using tensor methods. The algorithms that I have developed are not only of theoretical interest, but have also been impactful in practice. My methods are scalable to enormous datasets in social networks, document categorization and recommendation systems.

   I have developed a novel class of algorithms based on **tensor factorization of higher order moments of the observed data.** Tensors are higher order generalizations of matrices, and are essential for representing higher order relationships in data. I have characterized how higher order moments (typically third or fourth order) computed from data can be used to learn the parameters of various hidden variable models such as *Gaussian mixtures, latent Dirichlet allocation*, *hidden Markov models*, *network community models*, and so on [1–5]. These models are relevant in a wide range of applications in domains such as social network analysis, computational biology, document categorization, time series modeling, and so on. I have developed novel algorithms to manipulate and factorize the moment tensors to learn the model parameters. These algorithms retain the computational efficiency of many of the linear algebraic matrix method, and are embarrassingly parallel, which make them ideal for modern machine learning on large-scale data. **Thus, I have established that, for the first time, tensor algorithms can accurately learn the parameters of many latent variable models with polynomial computational and sample complexity.** The conditions for success are mild, and involve natural non-degeneracy conditions on the model parameters. **Tensor algorithms are more than two orders of magnitude faster than previous techniques for learning latent variable models (e.g. variational inference)**, while also achieving better accuracy of learning [6].

Neural networks have experienced a resurgence in machine learning, and have demonstrated tremendous performance gains in practice [7]. The popular backpropagation or stochastic gradient descent can however get stuck in bad local optima due to non-convexity of training neural networks, and therefore, have no provable guarantees. **For the first time, I have developed a novel method for training neural networks with guaranteed risk bounds with polynomial sample and computational complexity** [8]. I also demonstrate **how unsupervised learning can help in supervised tasks.** In this context, probabilistic score functions are estimated via unsupervised learning which are then employed for training neural networks under a tensor factorization framework. The conditions for success are mild: a small approximation error for the target function under the given class of neural networks, a generative input model with continuous distribution and general sigmoidal activations. My algorithm is based on tensor decomposition to estimate the network weights. Tensor methods come with guarantees of convergence to the globally optimal solution despite the non-convexity of the optimization problem. This results in first provable bounds for training a general class of neural networks.

The algorithms developed by me are not only of theoretical importance, but are also highly relevant in practice. To this end, I have proposed a number of algorithmic innovations to make tensor methods scalable. In my recent work [9], I have developed novel randomized sketching techniques based on efficient fast Fourier transform (FFT). This method has a running time which is independent of the tensor order, while a naive method has exponential time complexity. On another project, I am collaborating with researchers in NVidia to develop highly optimized libraries for tensor operations on both CPUs and GPUs based on extended BLAS kernels. BLAS or basic learning algebra subprograms are highly optimized linear algebraic libraries and these libraries are now being extended to tensor operations. In addition, my research group has developed tensor methods under the Spark platform, which has been very popular in a short amount of time. These innovations make it possible to build at scale tensor methods for learning in a range of domains such as social networks, recommender systems, text analysis, and so on, and I am striving to develop more open source packages on tensor methods.

4. Please attach one (or a few if you wish) graphic that best represents what your project is about. People are visual; one graphic can help them grab onto your project much quicker.

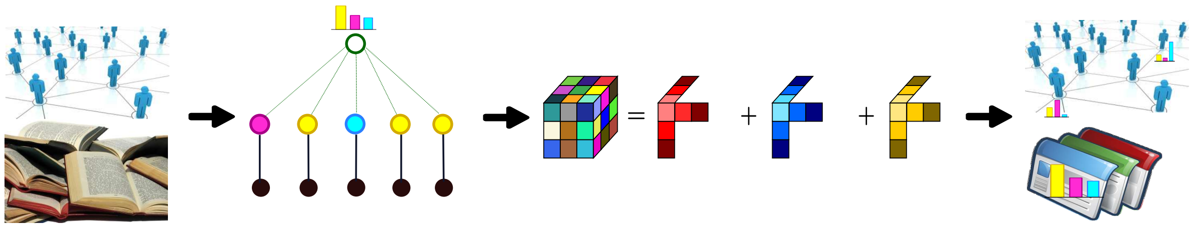

Figure 1: (i)Input: unlabeled data such as text and social network information. (ii)Posit probabilistic admixture models. (iii) Perform Tensor Decomposition. (iv) Use learnt parameters to carry out predictions.

2 Transitions

Describe anything from this project that you transitioned to anybody else (whether Army, DoD, govt, commercial, or other). a. Who did you give it to, and what is their organization? b. What did you give them? Code, papers, algorithms... c. What eventual application might this enable? d. What was your scientific accomplishment that enabled this?

Answer: We have released many open source software packages based on the tensor algorithms developed as part of this project. For example, we have a Spark package for tensor decomposition at `https://github.com/FurongHuang/SpectralLDA-TensorSpark`. We also have single machine implementation of tensor methods for topic modeling and community detection. More information can be found on PI Anandkumar's webpage.

We are also collaborating with a number of companies such as NVidia and Altera to develop more open source software, based on tensor algorithms, on a variety of platforms such as GPU and FPGAs.

There are plans to apply for SBIR/STTR grants based on the tensor algorithms which would enable more real world applications.

3 Awards/honors

For PI Anandkumar:

1. Early Career Excellence in Research Award, UCI Samueli School of Engineering. 2014-2015
2. Air Force Office of Sponsored Research (AFOSR) Young Investigator Award (YIP) 2015
3. Alfred P. Sloan Research Fellowship 2014
4. Microsoft Faculty Fellowship 2013
5. ARO Young Investigator Award (YIP) 2013
6. NSF CAREER Award 2013.

For graduate student: Furong Huang was awarded MLConf Industry Impact Award. She was one of the two students whose work, the committee believes, has the potential to disrupt the

industry in the future. See
`http://mlconf.com/mlconf-industry-impact-student-research-award-winners/`

References Cited

[1] A. Anandkumar, D. Hsu, and S.M. Kakade. A Method of Moments for Mixture Models and Hidden Markov Models. In *Proc. of Conf. on Learning Theory*, June 2012.

[2] A. Anandkumar, D. P. Foster, D. Hsu, S. M. Kakade, and Y. K. Liu. A Spectral Algorithm for Latent Dirichlet Allocation. In *Proc. of Neural Information Processing (NIPS)*, Dec. 2012.

[3] A. Anandkumar, R. Ge, D. Hsu, S. M. Kakade, and M. Telgarsky. Tensor Methods for Learning Latent Variable Models. *J. of Machine Learning Research*, 15:2773–2832, 2014.

[4] A. Anandkumar, R. Ge, D. Hsu, and S. M. Kakade. A Tensor Approach to Learning Mixed Membership Community Models. *J. of Machine Learning Research*, (15):2239–2312, June 2014.

[5] A. Anandkumar, D. Hsu, and A. Javanmard S. M. Kakade. Learning Topic Models and Latent Bayesian Networks Under Expansion Constraints. *Preprint. ArXiv:1209.5350*, Sept. 2012.

[6] F. Huang, U.N. Niranjan, M. Hakeem, and A. Anandkumar. Fast online tensor methods for learning latent variable models. *ArXiv 1309.0787, accepted to JMLR*, Sept. 2013.

[7] Yoshua Bengio, Ian J. Goodfellow, and Aaron Courville. Deep learning. Book in preparation for MIT Press, 2015.

[8] M. Janzamin, H. Sedghi, and A. Anandkumar. Beating the perils of non-convexity: Guaranteed training of neural networks using tensor methods. *arXiv preprint arXiv:1506.08473*, 2015.

[9] Yining Wang, Hsiao-Yu Tung, Alexander Smola, and Animashree Anandkumar. Fast and guaranteed tensor decomposition via sketching. In *Proc. of NIPS*, 2015.